

Die approbierte Originalversion dieser Dissertation ist an der Hauptbibliothek der Technischen Universität Wien aufgestellt (<http://www.ub.tuwien.ac.at>).

The approved original version of this thesis is available at the main library of the Vienna University of Technology (<http://www.ub.tuwien.ac.at/englweb/>).

Technische Universität Wien

Dissertation

Usability Issues in 3D Medical Visualization

ausgeführt

zum Zwecke der Erlangung des akademischen Grades
eines Doktors der technischen Wissenschaften

unter der Leitung von

Ao. Univ.-Prof. Dipl.-Ing. Dr. techn. Eduard Gröller,
Institut 186 für Computergraphik und Algorithmen,

eingereicht

an der Technischen Universität Wien,
Fakultät für Technische Naturwissenschaften und Informatik,

von


Dipl.-Ing. Andreas H. König,
Matrikelnummer 9225496,
A-7540 Neustift bei Güssing 64,
Österreich
geboren am 22. Juli 1974 in Graz.

Neustift, im April 2001.

Andreas H. König

Usability Issues in 3D Medical Visualization

(PhD thesis)



http://www.cg.tuwien.ac.at/~koenig/PhD_thesis
<http://www.cg.tuwien.ac.at/research/vis/vismed>
<mailto:koenig@cg.tuwien.ac.at>

The medical community has considered 3D visualization to be a toy since its beginnings. Although nice to behold, yet of no diagnostical value. Today, let us give researchers and companies in the field of visualization the chance to prove that we were wrong.

Dr. M. Prokop, MD, opening speech Spiral CT Workshop, Hannover, Germany, 1999.

We have seen the miracles and wonders of 3D visualization. The way, we are doing most of our work will change forever. Whoever is not convinced, shall wait. Wait, because these new tools have to be proven to be reliable. And it seems to me, that the usability still needs to be improved. State of the art approaches require us to be experts in an additional field: computer graphics. That has to change.

Dr. M. Prokop, MD, closing speech, Spiral CT Workshop, Hannover, Germany, 1999.

Abstract, Kurzfassung

(engl.) One of the most important technological developments in the 20th century is the use of computers in many areas of daily life, work, and research. Due to the new abilities in processing data, the amount of information that is handled, is much larger compared to previous times. As a consequence, the graphical representation, i.e., the *visualization* of data, established itself as a useful approach to investigate large data-sets. Among other application areas, the visualization of medical volume-data sets is a very important field of research. Anatomical or pathological structures acquired by medical imaging modalities like computed-tomography or magnetic-resonance scanners are visualized for further investigation. Many useful techniques have been developed in this area during the past decades. In this thesis the necessary adaptations and modifications to turn these scientific concepts into a software system for radiologists and other medical professionals are presented.

After a brief introduction to the visualization of medical volume data and a short overview over the state of the art, the main part of this thesis focusses on usability issues that arise when techniques developed by the scientific community of visualization experts are to be utilized in a commercial workstation. Additionally, three special topics of medical visualization are discussed, which were most sensible to usability issues: The specification of transfer functions, multimodal visualization of MR and fMRI data, and the handling of large-scale data sets as acquired by advanced multi-slice CT scanners.

(dt.) Eine der wichtigsten technologischen Entwicklungen im zwanzigsten Jahrhundert ist die Verwendung von Computern in vielen Bereichen des täglichen Lebens, der Arbeit und der Forschung. Aufgrund der neuen Möglichkeiten zur Datenverarbeitung ist die Menge an Information, die behandelt wird, viel größer als in früheren Zeiten. Als eine Folge davon hat sich die graphische Repräsentation, also die *Visualisierung* von Daten, als ein nützlicher Ansatz zur Untersuchung von großen Datensätzen etabliert. Die Visualisierung von medizinischen Daten macht einen sehr wichtigen Anwendungsbereich aus. Volumendatensätze, wie sie mit Hilfe von Computertomographie- oder Magnetresonanzscannern erzeugt werden können, dienen als die Grundlage für die Visualisierung von anatomischen oder pathologischen Strukturen des menschlichen Körpers. Viele verschiedene Techniken wurden zur Visualisierung von Volumendaten in den vergangenen Jahren entwickelt. In dieser Dissertation werden spezielle Probleme der Anwendbarkeit diskutiert, welche auftreten, wenn die aus dem akademis-

chen Bereich der Visualisierung stammenden Methoden für ein Softwaresystem verwendet werden, welches im täglich klinischen Betrieb zum Einsatz kommen soll.

Nach einer kurzen Einführung in den Bereich der Medizinischen Visualisierung und einem kurzen Überblick über den Stand der Wissenschaft, widmet sich der Hauptteil dieser Dissertation den Problemen, welche bei der Entwicklung eines Softwaresystems für den Einsatz im medizinischen Bereich auftreten. Danach werden drei spezielle Themen der medizinischen Visualisierung diskutiert, welche sich als im Hinblick auf ihre Anwendbarkeit als besonders herausfordernd gezeigt haben: Die Spezifikation von Transferfunktionen, die multimodale Visualisierung von MR und fMRI Daten und der Umgang mit besonders großen Datensätzen, wie sie von modernsten Mehrschicht-CT Geräten erzeugt werden.

Contents

Abstract, Kurzfassung	i
1 Introduction	1
1.1 Visualization	1
1.2 Medical Visualization	3
1.3 Evaluation	6
1.3.1 Usability	6
1.3.2 Reliability	6
1.3.3 Reproducibility	7
1.4 Overview	7
2 State of the art	9
2.1 Data acquisition	9
2.1.1 Computed Tomography (CT)	9
2.1.2 Magnetic Resonance Imaging	12
2.1.3 Functional scanning techniques	15
2.1.4 Other acquisition technologies	17
2.2 Volume Rendering	18
2.2.1 Ray casting	18
2.2.2 Shear Warp Factorization	19
2.2.3 Splatting	19
2.2.4 Maximum Intensity Projection	19
2.2.5 Hardware-accelerated rendering	20
2.3 Surface-based Rendering	21
2.3.1 Iso-surface extraction	21
2.3.2 Contour approaches	21
2.3.3 Segmentation masks	21
2.4 Segmentation	22
2.4.1 Region growing	22
2.4.2 Live wire / Intelligent scissors	22
2.4.3 Thresholding / Morphological Operations	22
2.4.4 Watershed Segmentation	22
2.4.5 Segmentation techniques tailored to blood vessels	23
2.5 Special Applications	23
2.5.1 Virtual Endoscopy	23
2.5.2 Surgical planning	24
2.5.3 Training	25
2.5.4 Presentation	25

3	Implementation: Tiani SpaceVision	26
4	Building a medical 3D workstation	27
4.1	Introduction	28
4.2	Hardware	29
4.3	Data Interface	31
4.4	Volume Reconstruction	32
4.5	Registration and Image Fusion	34
4.6	Segmentation	38
4.7	Reformatations	39
4.7.1	Multi-Planar Reconstruction	41
4.7.2	Curved-Planar Reconstruction	41
4.7.3	Reformatting the data set	43
4.7.4	Measurements	44
4.8	Rendering	45
4.9	Direct Volume Rendering	45
4.9.1	3D Windowing	46
4.9.2	Using transfer functions	47
4.9.3	Using indexed volumes	48
4.9.4	Tools	49
4.9.5	Maximum Intensity Projection	50
4.10	Surface rendering	51
4.11	Hybrid Rendering	52
4.12	High quality rendering	53
4.13	Virtual Reality devices	56
4.14	Storing the results	57
4.15	Conclusions	58
4.16	The future of 3D medical visualization	60
5	Multiple Views and Magic Mirrors	62
5.1	Introduction	62
5.2	The problems encountered in multimodal volume visualization	64
5.3	Index volumes for multimodal visualization	65
5.4	Multiple Views - Enhancing the visual impression of direct volume renderings	67
5.5	Magic Mirrors - Illustrating additional information	68
5.6	Results	69
5.7	Conclusions	70
6	Visualization of large data sets	71
6.1	Introduction	71
6.2	Requirements for the investigation of aluminium foam	72
6.3	Interactive Rendering	73
6.3.1	Threshold definition	73
6.3.2	Foam cell identification	74
6.3.3	Rendering acceleration	75
6.3.4	Cutting planes	76
6.4	Measurements	77
6.5	Focused Views	77
6.5.1	Cell size	77

6.5.2	Cell shape	77
6.6	Results	79
6.7	Conclusions	81
7	Advanced Transfer Function Specification	82
7.1	Introduction	82
7.2	Advanced transfer-function specification	83
7.3	Data range	84
7.4	Color range	89
7.5	Opacity range	91
7.6	Results	92
7.7	Conclusions	93
	Disclaimer	106
	Related URLs	107
	Curriculum Vitae	110
	Acknowledgements	114

Chapter 1

Introduction

1.1 Visualization

‘**visualize**’: to form a mental vision, image, or picture of (something not visible or present to the sight, or of an abstraction); to make visible to the mind or imagination.

Oxford English Dictionary, 2nd edition, 1989

At the turn of the millennium mankind was said to be enjoying the benefits of living in a so-called ”information society”. Shear endless amounts of information of arbitrary quality is available to us. With the rise of information technology it has become a difficult task to filter this information for the necessary essentials. Usually it is not feasible for humans to investigate the huge available amount of data by reading lists of numbers or other textual representations. The mapping of information into graphs or images, i.e., *visualization*, was already identified as a powerful tool for data investigation a long time ago. Leonardo da Vinci (1452–1519), for example, already used drawings to communicate scientific results. More recently, the extensive use of computers for data processing generated a new need for elaborated visualization techniques.

In the early 1990s annual-conference series, solely focusing on visualization, e.g., the ”Eurographics Workshop on Visualization in Scientific Computing” or the ”IEEE Conference on Visualization”, were established. Ten years later already a few compendia on visualization are available as comprehensive text books, for example, ”Scientific Visualization” by Gregory Nielson, Hans Hagen, and Heinrich Müller [NHM97].

To illustrate the role that visualization is playing at the beginning of the third millennium, some of the most important fields of application are listed below:

Medical data visualization – (anatomic) data is acquired through measurement devices, e.g., MRI or CT, which is then presented using volume visualization techniques, e.g., direct volume rendering or iso-surface extraction. This thesis will focus on the demands arising when scientific results in the field of medical visualization are applied for the construction of a software system, which is targeted to aid the daily clinical work of medical professionals.

Flow visualization – vector data, either computed by flow simulation, or measured data using experimental setups is plotted for the purpose of data investigation. For example, the design of new aircrafts can be checked using simulation and visualization without constructing expensive prototypes.

Geographic information systems (GIS) and visualization – for hundreds of years maps have been used as visualizations of geographic data. Techniques like color coding, height fields, iso-lines, and icons are used to show topographic information like mountains, rivers, etc. together with additional information, for example temperature.

Information visualization – big databases, multi-modal data, and abstract data (usually non-scientific data) increasingly require appropriate visualization techniques. Business data visualization (charts, diagrams, and graphs) is already widely used to illustrate economic data and relationships.

Visualization of microscopic data – molecules and/or atomic structures investigated in biology, chemistry, and physics are visualized for the task of analysis. Also data acquired by non-optical microscopes usually needs to be visualized before any investigation can be performed.

Large-scale data and visualization – astronomy, for instance, deals with data that is simulated or measured at a huge scale that prohibits direct investigation in most cases. Again, visualization can help to "fit", for example, entire galaxies into the study room of an astronomer.

Architectural visualization – planning of urban regions as well as buildings is simplified by the utilization of visualization methods. New buildings are visualized on the basis of computer aided design (CAD) data together with existing context. This allows to evaluate the feasibility of construction plans before actual construction. Additionally, the future buildings and their surroundings can be checked for obvious shortcoming during the planning stage.

Archeology and visualization – to investigate archaic cultures, for instance, visualization enables researchers to imagine life, habits, rites, etc. in former ages. Reconstruction of historic buildings using visualization is an area of increasing importance.

The following section discusses, how visualization is able to enhance the way medical professionals are able to diagnose, treat, and cure diseases and injuries of patients.

1.2 Medical Visualization

The radiology department is the heart of any hospital. Clinicians are able to detect epidemic diseases or problems with patients' metabolisms. But when it comes to the diagnosis of cancer, orthopedic irregularities, malignities of the brain, pulmologic problems, or subcutaneous injuries, the radiology department is the mother of all reports and findings.

Prof. W. Jaschke, University Hospital of Innsbruck, Austria, 1997.

The clinical discipline of radiology is providing 80% of all diagnoses for health problems of patients. Radiology departments in hospitals are the main source of data when it comes to the decision on the treatment of most kinds of diseases and injuries. An increasing number of stand-alone radiologists in suburban or provincial regions is providing services to all other kinds of medical professionals. Due to this immense importance, radiology has been the fastest growing field in the medical area for years.

At the turn of the 19th century *x-ray* technology was invented. Its diagnostic value was soon to be proven enormous. The discipline of radiology is constantly growing since this point in time. Since 1908 radiologist annually meet in Wiesbaden, Germany, at the *Deutscher Röntgenkongress* to share new insights and scientific breakthroughs. A number of other conferences and trade fares exist. The technology industry is very well aware of the fact that image acquisition devices and other tools for radiological use have become a billion dollar business. The market for these products has increased annually in average by 26% per year from 1990 to 2001, which makes it one of the fastest growing markets in the world.

Since the beginning days of radiology, doctors had to work in a very restricted way. Due to the radiation harm to the patient not more than about 3 x-ray images could be acquired for an examination. These images were made of sheets of film, which reacted to the exposure to the x-radiation. Radiologists would hang the developed films to light-boxes, which would give them enough brightness and contrast to examine the patients internal structures. With the rise of computed tomography (CT) more and more images could be acquired without increasing the radiation harm for the patient (see section 2.1.1 for further information on computed tomography). Additionally magnetic resonance imaging was introduced. Using this technology, even more images could be acquired, as no radiation harm is involved at all. Although CT and MR scanners are computerized machines, which are working with digital processing units inside, radiologists still used to print the images to sheets of film in order to hang them to the light-boxes. With the increasing number of images the light-boxes had to become bigger and bigger. Figure 1.1 shows a typical light-box filled with numerous MR images.

As the number of images, which could be acquired, further increased, the value of additional information on the patients anatomies became obvious. Solutions to the problems introduced with the increasing number of images had to be found:

- Light-boxes are unsuitable for a larger number of images. Even when so-called alternator systems were introduced, which would hold several dozens of films tacked to a band of plastic, scrolling the films into view



Figure 1.1: A traditional light-box, featuring a number of films showing MR images which are about to be diagnosed.

on the touch of a button, it was very cumbersome to manually hang the films to the light-box.

- Radiologic images have to be archived for at least ten years due to legal reasons. With the increasing number of images, the storage facilities in hospitals were soon not able to hold the films in an ordered and easily accessible way.
- The chemicals needed to develop the films represent a serious threat to our environment. Safe disposal or recycling were no longer options with the increasing number of films.

At about 1990 the computer industry started to take care of these problems. Picture Archiving and Communication Systems (PACS) were introduced. Images no longer had to be printed onto physical sheets of film. Image acquisition systems were connected to data storage servers by means of electronic communication networks. The patient's images are transferred as digital byte arrays obeying certain communication protocols, e.g. DICOM¹, which were soon introduced. Radiologists would use computer workstations to view and diagnose (radiologists prefer the term "to read the images") the images. Diagnostic workstations were developed, which could handle the large number of images and allow radiologist to view them in a convenient way. Figure 1.2 shows a diagnostic workstation, as it is used today.

Additional benefits arose with the introduction of PACS systems. All the workflow of a radiologist's job was easily integrated into the system: dictating the findings, converting the spoken findings to text by means of voice recognition systems, scheduling the appointments of examinations, automatic transfer of

¹<http://medical.nema.org/dicom.html>



Figure 1.2: A diagnostic workstation for "reading" all kinds of radiological images. The image shows a *Tiani JVision* solution.

findings and images to referring doctors, automatic billing for the services of the radiology department, etc.

When in 1999 the latest generation of CT scanners was presented, the dawn of medical 3D visualization as an applied technology had begun. Multi-slice scanners are able to acquire thousands of images within reasonable radiation and time limits. It became obvious, that even with the help of PACS and diagnostic workstations, it was impracticable to "read" on the basis of investigating single images. First of all, it is way too time consuming to investigate thousands of images one by one. Secondly, even a very experienced radiologist is easily overloaded with the amount of information contained in a multi-slice data set.

Reconstructing three-dimensional objects from the volume data sets is the answer to these problems. Scientific researchers have been dealing with medical volume visualization for several years. Their results always suffered from one problem: the resolution of the data sets was never high enough. With the introduction of multi-slice scanning, this problem was overcome and volume visualization could be taken to the highs of clinical application. Several companies started to build 3D-diagnostic workstations in order to enhance their PACS systems. Different problems arose when scientific research results of the academic community were to be used in a real-world application. This thesis is dealing with a variety of issues, which had to be solved in order to provide 3D medical visualization as a valuable tool for the daily clinical practice.

Four different ways of applying volume visualization to the medical field have been identified:

Diagnostics – Some new ways of diagnosis have been found. For radiologists, some diseases have become easier to identify and grade with the help of 3D visualization. For patients, some examinations have become less painful or uncomfortable. Some kinds of examinations were optimized in terms of time, so that a significantly increased number of patients can be treated per day.

Communication – Before the days of 3D medical visualization, three-dimensional representations of the patients internal structures were only present through the imaginal skills of expert radiologists. With visualization it is possible to communicate results to referring clinicians, surgeons, and patients by simply providing 3D renderings, which require no experience in interpreting slice images.

Planning – 3D visualization is able to aid the procedures of treatment, surgery, radiotherapy, orthopedics, etc.

Training – Students or young doctors can be trained for medical procedures of examinations or surgery using a virtual training environment. Using this new technology, more realistic training scenarios can be provided than with the traditional way of using dummy patients or body parts. Even the procedures which have to be employed in emergency situations can be trained this way. No patient has to undergo the risk of being harmed for training purposes.

All of these applications are going to be discussed in further detail in chapter 2.

1.3 Evaluation

As medical visualization is often dealing with the diagnosis and treatment of patient's diseases, highest demands on the quality of the utilized tools is necessary. Three aspects are important for guaranteeing the quality of results: Usability, reliability, and the reproducibility.

1.3.1 Usability

Medical software in general has to be highly optimized in term of usability. Every tool that is used in high-risk applications, e.g. power plant control or medical diagnosis, has to aid the users in preventing errors. For a medical workstation this means that the system should provide user guidance and assistance for all tasks of daily work.

1.3.2 Reliability

The user has to be assured, that non-obvious artifacts are not introduced into the visualization. Furthermore, no details must be omitted. All orientations (and positions in case of registered views) have to be preserved during visualization.

Future work might include the implementation of a concept like *ambiguity color coding*, which visually highlights regions in the visualization, where artifacts might be present (or details have been omitted) due to interpolation or sampling problems.

As reliability is the most critical aspect in the scope of quality assurance, proper testing in artificial scenarios as well as under real-world conditions has to be performed. The laws of most countries require software and hardware solutions applied to medical procedures to be approved by certain quality committees.

1.3.3 Reproducibility

Due to legal reasons all results obtained with a medical (3D) workstation have to be reproducible. In case of a wrong diagnosis it should be possible to track the source of the problem to the radiologist or to the company that developed the utilized tools.

Reproducibility is also important, when it comes to staging the progress of certain diseases. When a reconstruction that has been created some time ago is to be compared to a recent one, the same parameters for the visualization have to be used in order to guarantee, that obvious changes to the symptoms are not due to the changed nature of the visualization.

An important step to the reproducibility of results is the availability of the "session concept" mentioned in section 4.14. If all necessary steps in the generation of a visualization are saved, reproducibility is ensured.

1.4 Overview

The next chapter (No. 2, "State of the art") will describe the techniques which are available for the acquisition of medical volume data and the different computer-graphics techniques to visualize it.

Chapter 3 ("Implementation") briefly discusses, how the approaches presented in this thesis were integrated into a commercial software system.

Chapter 4 ("Building a medical 3D workstation") will cover all the usability issues, which have to be dealt with when trying to develop a software system for the daily practical use.

Chapter 5 ("Multiple Views and Magic Mirrors") will describe a solution we had to find for a specific field of application: multi-modal visualization.

In 1997 the introduction of the multi-slice scanning technology was announced. Until the first full-scale data sets were available (late 1998), we had to find a source of volume data, which would let us allow to prepare ourselves for the problems, which would arise with the huge medical data sets. In material science, industrial CT scanners are used to investigate the internal structure of workpieces. As there is no radiation harm, a high scanning resolution can be used, which results in large data sets that would be of approximately the same size as the expected multi-slice data sets. As solid pieces of metal are hard to compare to medical structures, we decided to use scans of a very novel field in material science: metal foams. These semi-hollow structures were comparable to the anatomical cavities of human patients to some degree. Chapter 6 ("Visualization of large data sets") describes the strategies, which had to be developed in order to be able to deal with large-scale data sets.

One of the most demanding problems in the area of medical volume visualization is the field of transfer-function specification. This field has been challenging for the scientific community since the beginning of research in the area of volume visualization. When it was eventually time to present the possibilities of visualization to medical doctors, they were forced to learn and understand the basics of volume visualization in order to be able to use the specification schemes used by computer-graphics experts. Chapter 7 ("Advanced Transfer Function Specification") presents an approach to the specification of transfer functions, which does not require any knowledge in the field of computer graphics. Even

the problematic specification of higher dimensional transfer functions is covered by this approach.

At the end of this thesis, a summary, conclusions, references to important literature, and acknowledgements will be given.

Chapter 2

State of the art

2.1 Data acquisition

The following section will describe the technologies, which are used to acquire volume data sets of the patients' anatomical structures.

2.1.1 Computed Tomography (CT)

A very important technology for the application of 3D visualization is computed tomography (CT). Scanned data values correspond to density values, which makes it perfectly suitable for the application of transfer functions, which are mapping density values into optical properties like color or transparency. The first derivative (the gradient) can be used to describe the homogeneity of regions. Regions with high gradient values describe transitions of material, e.g. bone to muscular tissue.

As CT was the first imaging technique capable of generating slice images, its development started quite some time ago:

1917: The basics for the mathematical calculus for the reconstruction of image slices is developed

1967: Godfrey N. Hounsfield concentrates on the technical realization of CT scanning

1971: A first prototype installation of a CT scanner in England is produced by EMI

1972: Serial production of the first skull CT "Mark I" is started

1979: The Nobel Price is awarded to Cormack and Hounsfield

The basic setup of a CT scanner is similar to a digital x-ray device. A x-ray tube is used to generate the radiation. A reference detector is used to track the amount of emitted radiation. The x-rays pass through the body of the patient and are attenuated by the different types of tissue. Eventually, another detector gathers the radiation, which passed through the patient. The measured value is the weakening w of the x-ray and is given by

$$w = \log \frac{I_1}{I_2}$$



Figure 2.1: A commercial CT scanner introduced in 1995.

where I_1 is the reference detector intensity and I_2 is the detector intensity. Using this basic setup, one-dimensional measurements of x-ray attenuation by the patient's body can be performed. By acquiring a multitude of such directions in a rotational way, the basic data for the reconstruction of a single slice is gathered.

Projection data may be acquired in one of several possible geometries described below, based on the scanning configuration, scanning motions, and detector arrangement. The evolution of these geometries is described in terms of "generations", and reflects the historical development. Current CT scanners are either third, fourth, or fifth generation machines, each having their own pros and cons.

- **First Generation: Parallel-Beam Geometry.** Parallel-beam geometry is the technically simplest and the easiest with which to understand the important CT principles. Multiple measurements of x-ray transmission are obtained using a single highly collimated x-ray pencil beam and detector. The beam is translated in a linear motion across the patient to obtain a projection profile. The source and detector are then rotated about the patient's iso-center by approximately 1 degree, and another projection profile is obtained. This translate-rotate scanning motion is repeated until the source and detector have been rotated by 180 degrees. The highly collimated beam provides excellent rejection of radiation scattered in the patient. However, the complex scanning motion results in long (approximately 5 minutes) scan times. This geometry was used by Hounsfield in his original experiments, but is not used in modern scanners.
- **Second Generation: Narrow Fan Beam, Multiple Detectors.** Scan times were reduced to approximately 30 seconds with the use of a narrow fan beam of x-rays and a linear detector array. A translate-rotate scanning

motion was still employed. However, a large rotation increment could be used, which resulted in shorter scan times. The reconstruction algorithms are slightly more complicated than those for first generation algorithms because they must handle narrow fan-beam projection-data.

- **Third Generation: Fan Beam, Rotating Detectors.** Third generation scanners were introduced in 1976. A fan beam of x-rays is rotated 360 degrees around the patient's iso-center. No translation motion is used. However, the fan beam must be wide enough to completely contain the patient. A curved detector array consisting of several hundred independent detectors is mechanically coupled to the x-ray source, and both rotate together. As a result, these rotate-only motions acquire projection data for a single image in as little as 1 second. Third generation designs have the advantage that thin tungsten septa can be placed between each detector in the array and focused on the x-ray source to reject scattered radiation.
- **Fourth Generation: Fan Beam, Fixed Detectors.** In a fourth generation scanner, the x-ray source and fan beam rotate about the patient's iso-center, while the detector array remains stationary. The detector array consists of 600 to 4800 (depending on the manufacturer) independent detectors in a circular arrangement that completely surrounds the patient. Scan times are similar to those of third generation scanners. The detectors are no longer coupled to the x-ray source and hence cannot make use of focused septa to reject scattered radiation. However, detectors are calibrated twice during each rotation of the x-ray source, providing a self-calibrating system. Third generation systems are calibrated only every few hours.

Two detector geometries are currently used for fourth generation systems:

1. A rotating x-ray source inside a fixed detector array
2. A rotating x-ray source outside a rotating detector array.

Both third and fourth generation systems are commercially available, and both have been clinically highly successful. Neither can be considered to be an overall superior design.

- **Fifth generation: Scanning Electron Beam.** Fifth generation scanners are unique in that the x-ray source becomes an integral part of the system design. The detector array remains stationary, while a high-energy electron beam is electronically swept along a semicircular tungsten strip anode. X-rays are produced at the point where the electron beam hits the anode, resulting in a source of x-rays that rotates around the patient with no moving parts. Projection data can be acquired in approximately 50 ms, which is fast enough to image the beating heart without significant motion artifacts.

An alternative fifth generation design, called the dynamic spatial reconstructor (DSR) scanner, is in use at the Mayo Clinic. This machine is a research prototype and is not available commercially. It consists of 14 x-ray tubes, scintillation screens, and video cameras. Volume CT images can be produced in as little as 10 ms.

The requirement for faster times, and in particular for fast multiple scans for three dimensional images, has resulted in the development of *spiral (helical) scanning* systems. Both third and fourth generation systems achieve this using self-lubricating slip ring technology to make the electrical connection with rotating components. This removes the need for power and signal cables which would otherwise have to be rewound between scans and allows for a continuous rotating motion of the x-ray fan beam. Multiple images are acquired while the patient is translated through the geometry in a smooth continuous motion rather than stopping for each image. Projection data for multiple images covering a volume of the patient can be acquired in a single breath-hold at rates of approximately one slice per second. The reconstruction algorithms are more sophisticated because they must accommodate the spiral or helical path traced by the x-ray sources around the patient. Figure 2.1 shows a Spiral CT modality.

Another breakthrough was made with the introduction of the *multi-slice technology*. Not only one ring of detectors is used, but multiple stacked rings. Therefore, with one revolution of the x-ray tube, not a single slice is acquired, but multiple ones. So in the same time, more data can be acquired without increasing the radiation dose, yielding data sets of higher resolution. The first models of multi-slice CT had four detector rings, the latest models (announced in 2001) feature 16 rings.

A numerical calculus is used to reconstruct density slices from the measured x-ray attenuations. Every pixel in a slice is an unknown variable, yielding a linear equation system. As more equations (attenuation values) than pixels in the slice are needed, the slice resolution is limited. The reconstruction is done by applying *Filtered Backprojection*, exploiting the correlation between Filtered Backprojection and Fourier reconstruction. Filtered Backprojection is used, because Simple Backprojection would just produce blurred images.

Please note, that due to the perpendicular arrangement of the patients body to the detector ring, images can only be acquired in axial direction.

Special CT techniques

One special field of computed tomography is *CT angiography (CTA)*. CT angiography is imaging the major arteries supplying any organ of the body. The exam is performed by obtaining rapid thin-section CT-scans through the artery or organ of interest during the rapid intravenous injection of contrast agent. Because of the rapidity with which the scanner obtains information, scan data is acquired when the arteries are highlighted by contrast. Using latest multi-slice technology, high resolution CTA scans can be acquired. The resulting volume data set can simultaneously demonstrate bony structures and their related vasculature. Especially for the diagnosis of diseases of the peripheral arteries of the legs or coronary arteries CTA is of interest.

2.1.2 Magnetic Resonance Imaging

Magnetic Resonance Imaging (MR or MRI) is the second imaging technique of importance to 3D visualization. Due to the technological basics of the scanning procedure, data values (also referred to as signal values) do not directly correspond to density values. Although the gradient of data sets can be very well exploited to identify structures, the application of transfer functions is not as

simple as for CT data. Segmentation techniques have to be used to separate individual structures for rendering.

In the following a brief history of MR technology is given:

1946: F. Bloch and E. M. Purcell discover the nucleus spin resonance

1971: R. Damadian identifies the possible use of nucleus spin resonance in medicine

1973: P. Lauterbur acquires the first MR image (of 2 water tubes)

1974: P. Lauterbur acquires the first image of a living being (a mouse)

1977: R. Damadian generates the first human MR scan (a whole body image which took 4.5 hours of acquisition time)

1980: the clinical use of MR technology booms

The theoretical basics of MR scanning are more sophisticated than for CT. An atom consists of a nucleus (proton) and a shell (electrons). The proton has a positive charge and rotates around its axis. The term "nucleus spin" refers to this action. Due to the spin the proton produces an electric current which induces a magnetic field. If an external magnetic field is applied, all nuclei align in the same direction, yet at different orientations, e.g. 1.000,000 protons might be at a high energy (oriented along the positive direction of the field), while 1.000,007 protons might be at a lower level (oriented towards the negative direction of the field). The protons move like a gyroscope, their movement is called *Precession*. The Larmor equation describes the frequency of precession ω_0 :

$$\omega_0 = \gamma B_0$$

where γ is the gyromagnetic ratio (42,5 MHz/T for protons), and B_0 is the magnetic power (given in Tesla). The precession movement yields a longitudinal magnetic vector. Inserting a high-frequency (HF) pulse at resonance frequency shortens the magnetic vector, and results in a transversal magnetic vector. After the pulse, the protons will take some time to realign to the magnetic field. This period is called *Relaxation time*. The longitudinal relaxation time is called $T1$ (usually in the range of 300 – 2000 msec). The transversal relaxation time is called $T2$ (usually in the range of 30 – 150 msec).

Fluids inside the patients body feature long $T1$ and $T2$ times. Fat features short $T1$ and $T2$ times. The high frequency pulse is repeated at a constant interval TR (Time to Repeat). TR can be used to highlight specific structures. A short TR emphasizes $T1$. If TR is very long, $T1$ has no influence any more and only a varying proton density influences the signal (more protons - higher signal) yielding a *proton density weighted image* (aka spin density weighted image).

Another scanning sequence is the so-called *Spin Echo Sequence*. Here a 180 degree pulse is inserted at $TE/2$ (TE = Time to Echo). The protons are back in phase at TE , but due to local magnet fields within tissues not all protons will be in phase.

For a *$T2$ Weighted Sequence* TE has to be chosen carefully: If TE is too small, the signal will have low contrast. If TE is too high, the image will have a bad signal to noise ratio. TE is usually chosen in the range of 30–80 msec.

The following settings are used for the most common scanning sequences:

Proton Density Weighted	Long TR, short TE
T2 Weighted	Long TR, long TE
T1 Weighted	Short TR, short TE

Figure 2.2 gives a visual comparison of the different scanning sequences.




Figure 2.2: The most common MR scanning sequences: T1 weighted (top left), T2 weighted (top right), and Proton Density weighted (bottom)

In order to highlight special structures (e.g., blood vessels), contrast agents are injected into the patient's vascular system. Contrast agents are "paramagnetic substances" with a local magnetic field and a shorter relaxation time. Contrast agents yield higher signal values in T_1 images and lower signal values in T_2 images. The most common contrast agent is *Gadolinium*.

High speed sequences are used to avoid motion artifacts. The following sequence types are most commonly used:

FLASH (Fast Low Angle Shot) – a flip angle for the high frequency pulse of 10 or 35 degrees (instead of 90 degrees) is used. This yields a faster recovery time for the protons.

GRASS (Gradient Recalled Acquisition at Steady State) – A gradient field is used, which induces more inhomogeneities yielding faster spin dephasing. The gradient field is flipped with a 180 pulse.

The trickiest thing about MR technology is the definition of a single slice. A gradient magnetic field ("Slice Selection Gradient Field") is used to specify a slice position (arbitrary directions of the gradient field are possible). A broader frequency spectrum of the HF pulse together with the steepness of the gradient field define the slice thickness. The signal is localized like this: After the HF

pulse a new gradient field is added: the "Frequency Coding Gradient Field", which changes the frequency of the protons within the rows. A short pulse of a third gradient field ("Phase Coding Gradient Field") accelerates the protons within the rows differently, which makes them change their phase. Eventually, frequency and phase are decoded to localize a single slice.

In order to apply the magnetic field, which aligns the protons, a magnet has to be used. Today magnets with a power rating between 0.5 and 1.5 Tesla are used. As a close to perfect homogeneity of the magnetic field is required, it is trimmed by "Shimming Magnets". Permanent magnets cannot be used in MR scanners, as they are thermally instable and too heavy. A permanent magnet capable of producing a field of 0.3 Tesla would weigh 100 tons. Standard electro magnets cannot be used either as they get too hot. Supra conducting magnets are the technology of choice. Cryogen is used to cool the magnet to -269 degrees Celsius. If the temperature raises, a "quench" valve is used for stabilization.

Different coils are used inside the MR scanner:

Volume Coil – sends HF pulse and measures response from big body parts

Skull Coil – used for brain measurements

Shim Coils – used for trimming the homogeneity

Gradient Coils – for all three axes

Surface Coils – used to measure tissues close to the surface - HF pulse is sent by the volume coil

Further information on MR technology can be found in specialized literature [CS92].

Special MR techniques

An interesting variant of classical MR scanning is *MR angiography*. Its application is the investigation of vascular structures. Combining a well timed gadolinium contrast agent injection with a three-dimensional (3D) spoiled gradient-echo (SPGR) MR sequence yields a good signal-to-noise ratio. Using this fast acquisition sequence, MR angiography is capable of covering extensive regions of vascular anatomy within a single breath-hold of the patient.

2.1.3 Functional scanning techniques

Functional MRI

Functional MRI (fMRI) is a technique for determining which parts of the brain are activated by different types of physical sensation or activity, such as sight, sound, or the movement of a subject's fingers. This "brain mapping" is achieved by setting up an advanced MRI scanner in a special way so that the increased blood flow to the activated areas of the brain shows up on functional MRI scans.

The subject in a typical experiment will lie in the MR scanner and a particular form of stimulation will be set up. For example, the subject may wear special glasses so that pictures can be shown during the experiment. Then, MRI images of the subject's brain are taken. First of all, a high resolution single scan is acquired. This is used later as a background for highlighting the brain areas

which were activated by the stimulus. Next, a series of low resolution scans are acquired over a period of time, for example, 150 scans, one every 5 seconds. For some of these scans, the stimulus (in this case the moving picture) will be present, and for some of the scans, the stimulus will be absent. The low resolution brain images in the two cases can be compared in order to see which parts of the brain were activated by the stimulus.

Two dominant tissue contrast mechanisms feature functional sensitivity in MR imaging and are produced via haemodynamic responses. Precise changes in brain activation or metabolism are not directly observed, but the effects of local increases in blood flow and microvascular oxygenation on one or more of the above mentioned MR mechanisms can be mapped as a change in raw image intensity.

One mechanism depends upon the fact that the microvascular MR signal on T_2 and T_2^* weighted images is strongly influenced by the oxygenation state of the blood. The rate of loss of proton-spin phase-coherence is a measure of T_2 and local magnetic field homogeneity (T_2^*). This can be modulated by the presence of intravoxel deoxyhaemoglobin. Research has shown that the observed T_2^* is dependent on the presence of blood deoxygenation and that deoxygenated haemoglobin is a *blood oxygenation level dependent* or *BOLD* effect that can be observed by noninvasive MR imaging at high magnetic fields.

The BOLD imaging technique does not measure tissue perfusion or flow directly, however, because over 70% of the brain's blood lies within the microvascular capillaries and venules. The measurement of the magnetic susceptibility-induced T_2^* signal loss is thought to most reflect the regional deoxygenation state of the venous system. In addition, proton perfusion and diffusion through changing local gradients modulated by changing oxy-/deoxyhaemoglobin levels has a direct impact on the observed T_2 relaxation times, which is another mechanism of tissue contrast generation. Amongst these various mechanisms, the T_2^* effect is more significant by factors of 3 to 10 and is the dominant and most widely-studied mechanism employed in fMRI.

In short, the response to a local increase in metabolic rate is increased delivery of blood to the activated region. Such a change in haemodynamics produces small alterations in T_1 , T_2 or T_2^* , which can be visualized as a change in MR image intensity.

The functional data acquired at fMRI examinations has to be displayed in conjunction with the underlying anatomical basis. This special field of multi-modal visualization is discussed in chapter 5.

Positron Emission Tomography

Positron Emission Tomography (PET) is an imaging technique that uses a radioactive tracer to show chemical activity of anatomical structures. PET is particularly useful for the detection of cancer, coronary artery disease and brain disease. The PET scanner for instance pinpoints the destination of radioactively tagged glucose, oxygen, or drugs to reveal the parts of the brain involved in performing an experimental task.

Like CT and MRI, PET relies on computerized reconstruction procedures to produce tomographic images. It is performed by means of detecting positron emission by use of tomography. Two ways in which radionuclides decay that will reduce excess positive charge on the nucleus include the neutralization of

a positive charge with the negative charge of an electron or the emission of a positron from the nucleus. The positron will then combine with an electron from the surroundings and annihilate. Upon annihilation both the positron and the electron are then converted to electromagnetic radiation. This electromagnetic radiation is in the form of two high-energy photons which are emitted 180 degrees away from each other. It is this annihilation radiation that can be detected externally and is used to measure both the quantity and the location of the positron emitter. Simultaneous detection of two of these photons by detectors on opposite sides of an object places the site of the annihilation on or about a line connecting the centers of the two detectors. At this point mapping the distribution of annihilations by computers is allowed. If the annihilation originates outside the volume between the two detectors, only one of the photons can be detected, and since the detection of a single photon does not satisfy the coincidence condition, the event is rejected. Simultaneous detection provides a precise field of view with uniform sensitivity. This occurs because wherever the disintegration takes place between the two detectors, the photons must in sum have travelled the full inter-detector distance in order that the event can be recorded.

Single Photon Emission Computed Tomography

In perfusion *single photon emission computed tomography* (SPECT), image intensity is directly related to cerebral perfusion or blood flow, which in turn relates directly to brain functional activity.

As in CT, SPECT imaging involves the rotation of a photon detector array around the body to acquire data from multiple angles. Using this technique, the position and concentration of radionuclide distribution is tracked.

SPECT imaging is inferior to PET because of the attainable resolution and sensitivity. Different radionuclides are used for SPECT imaging that emit a single photon (usually about 140keV), rather than positron emission (resulting in two high energy 511 keV photons) as in PET. Because only a single photon is emitted from the radionuclides used for SPECT, a special lens known as a collimator is used to acquire the image data from multiple views around the body. The use of a collimator results in a tremendous decrease in detection efficiency as compared to PET, because only 1 or 3 collimators are used. Higher resolutions require better sensitivity and the resolution of SPECT is way less than that of PET. The resulting useful resolution (about 7mm) for SPECT is inferior to PET resolutions by a factor of 3 or 4. Although SPECT imaging resolution is not that of PET, the availability of new SPECT radiopharmaceuticals, particularly for the brain and head, and the practical and economic aspects of SPECT instrumentation make this mode of emission tomography attractive for clinical studies of the brain. Also, the cost of SPECT imaging is about one third of the cost of PET imaging.

2.1.4 Other acquisition technologies

3D Sonography

One of the latest imaging technologies capable of acquiring volumetric data is *3D sonography*. Like for conventional ultrasound examinations a probe is used,

which scans the internal structures of a patient by means of sound waves. By moving the probe, the continuously scanned information can be integrated to a data volume. As it is nearly impossible to manually move the probe accurately enough for this task, it has been automated. Volumetric probes perform an automatic scan through a mechanical movement of the head of the probe so that very close planes are integrated into a volume. Time for completing a scan is from 3 to 10 seconds, depending on the required quality. The acquired volume typically has the shape of a pyramid.

In general it has to be stated, that the quality of 3D-sonographic data-sets does not necessarily support the application of 3D visualization for sophisticated diagnostic purposes due to inherent noise artifacts. Nevertheless, it can be applied in order to show a mother-to-be a 3D reconstruction of her baby, proving, that it does not suffer from any deformational diseases.

2.2 Volume Rendering

Volume rendering is concerned with the display of volumetric data sets in their entirety. As surface-based techniques (see section 2.3) focus on emphasizing material transitions, volume rendering is able to display all kinds of data configurations. Using color and transparency classifications, even homogeneous regions can be rendered as cloud-like structures. Figure 2.3 gives a visual comparison.




Figure 2.3: Surface Rendering (left) vs. Volume Rendering (right) of the same data set.

A variety of different approaches have been developed up to now. The most important techniques will be briefly discussed in the following sections.

2.2.1 Ray casting

Ray casting was one of the first volume rendering techniques introduced in 1988 by Mark Levoy [Lev88]. The *ray tracing* method, which is common in the field of photo-realistic rendering is adapted to the need of volume data. Originating from the camera position, viewing rays are traced through the data volume. As optical properties are evaluated for every encountered data voxel [Max95, DCH88], the color of every pixel of the final image can be calculated by means

of integration. Sampling the data volume along the viewing rays at equidistant locations is computationally expensive, as it requires interpolation techniques to find data values in-between voxel positions. Several acceleration techniques have been proposed in order to speed up the ray casting process [AHH⁺94, GHY98]. Some approaches are exploiting the fact, that there are regions in the volume data set that do not contribute to the final image [YS93, SK00]. Other approaches try to optimize the way viewing rays are sampled, either by taking the final image into account [WKB99, KBD⁺98], or by using adaptive sampling strategies [OOLP98]. Also the way, viewing rays are traversed through the data volume are subject to acceleration approaches [LY96, YK92]. Finally the coherency between adjacent viewing rays or frames of an animation is used to speed up the rendering process [BJNN97]. Special acceleration techniques have been developed for ray casting in the field of Virtual Endoscopy [HO00, YHW⁺97] and segmentation [TD97].

Although these acceleration techniques yield a significant speedup in rendering times, ray casting (without a tradeoff in image quality) is far from being an interactive technique. Nevertheless, it generates images of highest quality.

2.2.2 Shear Warp Factorization

The *Shear Warp Factorization* method was introduced by Philippe Lacroute and Marc Levoy in 1994 [LL94]. A significant gain in rendering performance is achieved by rearranging the voxels in the computer's memory space to allow an optimized ray traversal (shearing of the volume data set). As this procedure yields a distorted image, a final correction step is necessary (warping the image). Shear warp rendering is capable of generating images of reasonably sized data volumes at interactive frame rates. Nevertheless a good deal of image quality is sacrificed. For this reason several acceleration techniques have been developed, using parallel rendering approaches [Lac96], or optimized sparse data representations [Csé99]. Additionally techniques for the hybrid rendering of volume data and polygonal structures based on shear warp factorization were introduced [ZYS99].

The reasonable speed/quality tradeoff of this technique made it the method of choice for implementation in special hardware architectures (refer to section 2.2.5).

2.2.3 Splatting

The *splatting* method was introduced by Lee Westover in 1990 [Wes90, Wes91]. Volume rendering displays are composed by accumulating "splatted" footprints of data voxels on the image plane. This object space method showed to be very well optimizable by using sparse data sets and other strategies [HMSC00, MY96, OM00, LH91]. Problems with aliasing artifacts were the topic of advanced research [MMS⁺98, SMM⁺97] as well as the improvement of the visual quality of images [MMC99].

2.2.4 Maximum Intensity Projection

Maximum Intensity Projection (MIP) is the rendering technique of choice when it comes to the investigation of vascular structures. Unlike the accumulation

process of other volume rendering techniques, the data volume is not integrated along the viewing direction. Only the maximal data value encountered along a viewing ray is contributing to the final image. Nevertheless, its computational effort is considerable. Acceleration techniques, which work well for ray casting, like *early ray termination* or *space leaping* cannot be employed. For this reason several other acceleration techniques were developed for MIP [SGS95, CS98, MKG00, MHG00, MGK99, HMS95]. Further enhancements include approaches to improve the three-dimensional impression of MIP renderings, like *Depth-shaded MIP* or *Local MIP*. Figure 2.4 gives a visual comparison of the different techniques.




Figure 2.4: Classic Maximum Intensity Projection (MIP) (top left), Local MIP (top right), and Depth-shaded MIP (bottom)

2.2.5 Hardware-accelerated rendering

In recent years hardware accelerated rendering systems became available. Classical volume rendering techniques are implemented to run on dedicated high performance signal processors. Implemented techniques include *Volume Rendering using 3D textures* [MHS99], *Parallel Volume Rendering* [RPSC99], and *Shear Warp Factorization* [PHK⁺99].

Hardware-based rendering acceleration seems to be a feasible bridge solution for the time being, but as all-purpose computers constantly increase in power, such dedicated solutions might not be needed any more in the future. Furthermore, the high performance components used for accelerators are always custom tailored to the solutions. For this reason, the accelerators are expensive, have

long development times, and are trying to avoid being overpowered by software-only solutions on PC-based systems, which evolve at a much faster rate than the dedicated hardware solution. Today, the hardware accelerators are still a little ahead in terms of rendering performance, but that will change.

2.3 Surface-based Rendering

Surface Rendering or *Surface Shaded Display* (SSD) is dealing with the display of surfaces extracted from volume data sets. Three different approaches for the extraction of surfaces have been developed: iso-surfacing techniques, contour connection techniques, and the display of segmentation results (so-called "masks").

2.3.1 Iso-surface extraction

The *Marching Cubes* technique is the best established iso-surface extraction approach, introduced by W.E. Lorensen and H.E. Cline in 1987 [LC87]. The original technique suffered from the problem of ambiguities in the surface construction scheme, which would result in holes in the extracted surfaces. Further research work corrected this problem [NH91, Mat94, Sho98, MSS94b]. The decimation of high-resolution surfaces generated by the Marching Cubes technique was another problem [SFYC96]. The time consuming process of the surface extraction motivated researchers to come up with acceleration strategies [MSS94a, WG92, LSJ96, IK94, LA98, DMMN99]. For special applications like Virtual Endoscopy a view dependent extraction technique was proposed [LH98]. Also in the field of network-based visualization recent research work has improved the original Marching Cubes technique [EWE99].

2.3.2 Contour approaches

Before the introduction of iso-surfacing techniques, contour connection approaches were widely used. In a semiautomatic or manual way the contours of structures of interest were identified on single slices. Geometrical algorithms were then used to connect the contours of adjacent slices in order to form a closed surface [AN94, MSS92, KSS98, Mey94]. Enhancements to this technique include the generation of meshes with more or less uniformly shaped triangles, which could be reduced in size (and detail) by standard mesh reduction algorithms [FKU77]. The major problem of this kind of surfacing technique, which took some years to be solved by science, was the handling of branches in the objects' structures [Sha81].

Contour connecting approaches are no longer as widely used as they used to be.

2.3.3 Segmentation masks

A natural application for surface-based rendering is the display of regions of a data volume, which have been identified by segmentation techniques. As such segmentation masks are binary volumes of arbitrary structure, the straightforward application of iso-surfacing techniques yields staircase artifacts, which

spoil the visual expression of any rendering. For this reason, special smoothing techniques are employed to get visually satisfying renderings of segmentation masks [WE97, Tau95, Neu00, Gib98, LG99].

2.4 Segmentation

The field of segmentation has been a comprehensive area of research for decades. Disciplines like *Computer Vision*, *Pattern Recognition*, and *Medical Visualization* often depend on segmentation results. In order to keep the scope of this thesis reasonably focussed, only techniques that can be successfully applied to tasks of visualization are mentioned in the following sections.

2.4.1 Region growing

Region Growing is a statistical segmentation technique. An initial starting region is enlarged, as long as the added voxels fulfill certain connectivity criteria. These criteria can be defined to include the data value, several derivations of the data value, texture information, etc [XAW95, RJ97]. This technique is yielding excellent results, but it is computationally very expensive. Parallelizing can help to speed up the calculation [CRFS94].

2.4.2 Live wire / Intelligent scissors

The *Live Wire* or *Intelligent Scissors* approaches work essentially in the same way [BM97, MB95]: The user is tracing a desired contour by moving a pointer vaguely along the transition of different tissues on an image slice. The algorithms are using the gradient information of the volume data to correct the users movements according to the underlying material transitions. With the introduction of a faster version of Live Wire [FU00b], this technique has become the semiautomatic segmentation tool of choice. An extension to the third dimension has also been proposed [FU00a].

2.4.3 Thresholding / Morphological Operations

Thresholding is one of the simplest segmentation tools. The user just defines a range of data values, and all voxels in the volume-data set within this range are marked to belong to the segmentation mask. The drawback is obvious: If two connected structures in this range are to be separated, special tools have to be employed. *Morphological Operations* are filter operations working on the volume data set, which are able to break thin connections, fill or open holes, and to perform a variety of other tasks [HH92, SBTH92]. Although this set of tools is a very powerful one, a lot of interaction by experienced users is required. Additionally, processing times for the filtering of the volume-data set can be enormous. For these reasons, Thresholding and Morphological Operations are challenging to be used in daily medical practice.

2.4.4 Watershed Segmentation

Although *Watershed Segmentation* has been introduced to the field of image processing a long time ago [BL77], it had to be revived for the use in medical

applications [Beu91]. In this semiautomatic technique, the user is restricting the desired segmentation result by defining a set of points, which either do belong to the structure of interest, or do not belong to it. Using the physical analogy of a rising water level, the gradient of the image data is interpreted as ridge and valley lines. When the water level increases, pools separated by lower gradient transitions will be merged earlier than others, which are separated by high gradient transitions. Only two pools of "water", where one of them contains voxels pre-specified as "in the desired region" and the other one contains voxels "outside the desired region", are not merged. When the "water-level" has reached its maximum height, the two types of pools identify background and segmentation mask.

With a real-time implementation of this technique it is possible to define and refine a segmentation with just a few clicks of the mouse. Although developed for 2D images, an extension to volumes is straightforward, making this approach the technique of choice for semiautomatic segmentation.

2.4.5 Segmentation techniques tailored to blood vessels

Some mentionable techniques have been introduced for the segmentation of blood vessels. Due to the well defined geometry of blood vessels and other tubular structures, these algorithms are able to work automatically, requiring only minimal user interaction.

Vessel Tracking is an algorithm, which is able to not only generate a segmentation mask of the desired vessel, but also a topological description of the vessel's branching structure [WNV00]. Similar results can be achieved with the so-called *Wave* algorithm [Göp97]. Another segmentation technique capable of automatically segmenting tubular structures is *Multi-scale 3D-filtering based segmentation* [FNVV98, SNS98, SWBN00].

2.5 Special Applications

In the following sections some specialized applications of 3D visualization are discussed.

2.5.1 Virtual Endoscopy

Virtual Endoscopy is an application which deals with the exploration of hollow organs and anatomical cavities using volume data. Using a real endoscope is always painful or at least inconvenient for the patient. Virtual Endoscopy is used as a non-invasive diagnostic technique. It may also be applied in surgical planning, training and exploration of structures not reachable with a real endoscope. All these fields have similar requirements concerning the visualization system: accuracy, intuitive interaction, fast visualization, and a short preprocessing time.

The common approach for the visualization of a virtual endoscopy is surface rendering, yielding images close to a real endoscopy. If external structures are of interest, volume rendering techniques have to be used. These methods do not display the exact shape of the inner lumen very well. For certain applications, e.g., surgery planning of a transbronchial biopsy, both, the shape of the inner



Figure 2.5: Virtual Angioscopy of the aorta. Top left: surface-based overview rendering; Top right: volume rendering overview; Bottom left: reference slice image; Bottom right: endoluminal view.

lumen as well as outer structures like blood vessels and the tumor have to be delineated. A new method has been developed, that allows a quick and easy hybrid visualization using overlays of different visualization methods like different surfaces or volume renderings with different transfer functions in real time on a low-end PC [WVH⁺00]. To achieve real time frame rates, image based rendering techniques have been used. Please refer to figure 4.8 for an image of a Virtual Bronchoscopy. Another example shows a Virtual Angioscopy (see figure 2.5).

Virtual Endoscopy is mainly employed for the investigation of the colon (Virtual Colonoscopy), of blood vessels (Virtual Angioscopy), and of the trachea (Virtual Bronchoscopy).

2.5.2 Surgical planning

Another important field of application for a medical 3D workstation is the planning of surgical procedures. Especially when it comes to the task of transplanting kidneys or livers, the configuration of essential blood vessels can be different for every patient. As transplantation procedures are very time consuming and involve a certain risk for donor and recipient, it is desirable to shorten the necessary amount of time by proper planning. Before the procedure, radiologists and surgeons team-up in order to interpret the available radiological information and data. By using 3D reconstructions, the radiologists are able to brief the surgeons with the actual situation inside the patient. Potentially dangerous hazards can be identified and highlighted in the 3D visualization, e.g., for surgical procedures involving blood vessels, it is important to know anatomic variants of vascular anatomy. Other examples can be found in the field of neu-

rosurgery. When a surgeon has to treat a tumor inside the brain, it takes proper planning to avoid major blood vessels, when opening a channel from the skull to the actual location of the tumor. A medical 3D workstation provides the tools to find the optimal entry point by interactively examining a reconstruction of the brain, the skull, and the vascular system. During this "briefing" session, the surgeon is even able to perform a "test run" of the procedure and identify potential hazards in order to modify his strategy accordingly.

During the medical procedure 3D visualization can provide valuable guidance by utilizing an augmented reality system (refer to section 4.13). Although these systems are not widely established yet, most surgeons at least rely on still images of 3D reconstructions. Some surgeons are already starting to use an interactive display of a 3D workstation in the surgery theater.

2.5.3 Training

3D visualization can also be used as a valuable tool for the training of students. Various fields of application are possible:

Anatomy – For centuries students of the medical disciplines had to study anatomical structures from 2D still images. Dissection courses have to be used to enable the students to experience the real 3D shape of structures they only know from the atlases. With the help of a medical 3D workstation, real world data sets can be used to generate interactive models of the structures of interest, enabling the students to gain a better understanding of anatomy.

Endoscopy – The training of student endoscopists involves certain risks for real patients. Without a lot of experience it is nearly impossible to direct the movements of a real endoscope without harming the patient. For this reason endoscopists have to undergo long training periods using "dummy" patients, which only prepares them partly for their first real patient. A valuable add-on to this training is the utilization of a virtual training scenario. Using a special input device with haptic feedback, which emulates the behavior of a real endoscope, connected to a 3D workstation, the endoscopy procedure can be trained. Collision-detection algorithms can be used to provide real-time feedback to the student, concerning the movements of the virtual tools.

Surgery – Special surgical procedures, like modifications of the visual nerves, require special experience, which cannot be gained on real patients. Virtual training scenarios with tracked surgical instruments can be used to teach these procedures.

2.5.4 Presentation

Last, but not least, the presentation of medical cases is an application for any medical 3D workstation. It is easier for the radiologist to communicate the basis for his finding to clinicians, surgeons, and even patients by using a reconstructed visualization or even an interactive display. For the presentation of interesting special cases in printed media, high resolution renderings can be generated (refer to section 4.12).

Chapter 3

Implementation: Tiani SpaceVision

The concepts described in this thesis were used to develop a commercially available 3D visualization workstation. This was done at Tiani Medgraph AG, Brunn am Gebirge, Austria. The workstation is called JVision/SpaceVision and is sold for about 100.000,- USD.

A development team of 6 people was involved in the implementation process. It was started in the middle of 1997 and is still going on today.

Java was the programming language of choice, making SpaceVision web enabled and platform independent. Hardware rendering acceleration (VolumePro technology by Real Time Visualization Inc.) is used along with software based rendering techniques in order to provide realtime rendering feedback to the user.

SpaceVision stations are in use in hospitals like the University Hospital of Innsbruck, Austria; Hospital Rudolfstiftung, Vienna, Austria; County Hospital Kirchdorf, Austria; County Hospital Steyr, Austria; AKH (general hospital) Vienna, Austria; Marienhospital Stuttgart, Germany; County Hospital Murnau, Germany; the Röhkliniken (12 hospitals), Germany; and the West Coast Radiology Center, Santa Ana, USA. The number of customers interested in buying and using a SpaceVision station is rapidly increasing.

Chapter 4

Building a medical 3D workstation

This chapter is targeted on fulfilling three tasks:

1. It defines the requirements for a medical 3D workstation.
2. It describes how it was tried to meet these requirements during the development of the Tiani SpaceVision workstation.
3. It gives a comparison on how other workstations are dealing with these requirements.

Five of the most established state-of-the-art workstations were chosen for the comparison:

- **JVision/SpaceVision 1.8** by *Tiani Medgraph*, Austria, www.tiani.com.
- **Virtuoso** by *Siemens Medical Systems*, USA, www.siemensmedical.com.
- **Advantage Workstation 4.0** by *GE Medical Systems*, USA, www.gemedicalsystemseurope.com.
- **Vitrea 2** by *Vital Images*, USA, www.vitalimages.com.
- **Plug'n'View 3D v.2.7** by *Voxar Ltd.*, UK, www.voxar.com.

The descriptions and judgements on these workstations given in this thesis are based on information provided by representants of the companies that develop the workstations, on information provided via the internet, and on hands-on experience by the author. Some of the evaluated criteria (e.g., usability) are somewhat subjective. The author has tried to find a more general judgment by interviewing radiologists on these properties.

Errors or omitted information in the description of the workstations are of course unintentional. The comparison has been prepared to the best knowledge of the author and with all information available to the author. The author shall not be held liable for any misunderstandings or misinformation in this publication.

	CT	MR
Type	CT Angiography	MR T1 weighted
Body region	hip to toes	entire head
Intra-slice resolution	512x512	256x256
Number of slices	1200	400
Slice distance	1 mm	0.75 mm
Signal resolution	12 bit	8 bit
Size of data set	630 MB	39 MB
Scanning time	1 min 12 sec	4 min 14 sec

Table 4.1: Properties of example data sets

4.1 Introduction

For purposes of medical diagnosis it is most desirable to have at hand as much detailed information as possible on the anatomical and pathological features of the patient. For this reason, examinations by the various scanning modalities in clinical practice try to find the optimal relation between scanning resolution and harm or inconvenience for the patient. With the recent improvement of modalities featuring high spatial scanning resolutions, special strategies for the investigation of huge data sets are more necessary than ever. Especially Spiral-CT or Multi-Detector CT scanners are capable of acquiring several hundreds of CT slices within a reasonable amount of time and with minimal radiation harm for the patient. The same is true for MR modalities. With MR the spatial resolution within a slice is usually lower, but there is no radiation harm for the patient at all. Therefore, the size of the acquired data sets is merely limited by the time, the patient is supposed to stay inside the MR scanner. Table 4.1 states the most important properties of example CT and MR data sets.

Traditionally, a radiologist supposed to diagnose such data sets would look at the scanned images slice by slice. By mentally reconstructing the sampled information into a three-dimensional representation, he would judge on the health of the patient. This procedure requires the radiologist to have well-founded experience as well as a highly sophisticated understanding of human anatomy. To create a "complete" mental image of the patients structures, the radiologist has to take all available information into account. It is obvious, that looking for a single patient at hundreds of slices (and soon thousands!) is way too time consuming.

Although these new modalities provide reliable enough information to find even small (and still curable!) pathological structures, it has become impossible to do diagnosis the traditional way. 3D medical visualization can help to overcome this problem.

A future vision of radiological practice could be like this:

- Patients attend *screening examinations* once a year. As scanning and evaluation by the radiologist takes just a couple of minutes, it is affordable by means of time and money.
- A high resolution data set of the patient's entire body is acquired.
- The radiologist is presented a hierarchical representation of the data: He is checking for possibly pathological structures in a comprehensive 3D

context of the patients body. When interesting regions are identified, the radiologist does further investigations by *focussing* on a more detailed representation of this region.

- When surgical intervention for the diagnosed problem is necessary, the radiologist prepares an information suite, which aids the surgeon in identifying crucial structures during the operation by utilizing *Augmented Reality*¹ techniques.

3D medical visualization is not only able to do the reconstruction of the 2D slices into a 3D representation. It also may provide the user with displays incorporating different levels of information integration. Like in the example mentioned above, the user could start his investigation on a very coarse level of the data representation and gradually refine the representation in special regions in order to avoid information overload and focus on the important parts. Additionally different properties of the patient's data important for the diagnosis can be presented to the radiologist, e.g., curvature information (which is not simply derived in a visual way) can be color coded, selective displays can hide or emphasize certain parts or structures, etc.

The features to be integrated into a medical workstation have to be carefully chosen from a large number of possibilities. On the other hand, the requirements and special demands by radiologists (customers, in our case) are challenging. The following sections will discuss the requirements vs. different ways to fulfill them, as well as new possibilities in diagnosis, which have shown up during the development of our workstation. The topics of hardware, data integration, volume reconstruction, registration, segmentation, reformatations, different rendering techniques, special applications, and post processing will be covered. Eventually, detailed information on the evaluation of our system will be given. Conclusions and some visions for the future will conclude this chapter.

4.2 Hardware

A very important decision to be made when starting to build a visualization software system is the choice of the hardware platform to be used. The market for computer hardware is constantly changing and evolving. A very good example for demonstrating the highly dynamic developments in this field are *OpenGL* graphics adapters. Up to one year ago, only high-end graphics adapters could deliver the performance to meet the enormous demands of medical visualization. Nowadays even cheap consumer-graphics adapters are almost powerful enough. The gaming industry is pushing the development of increasingly faster and cheaper graphics adapters.

It is tempting to use such units for a medical workstation, as these adapters are cheap, sold in very large quantities, and therefore tested and developed very thoroughly. Nevertheless, it has to be seen, that the performance features,

¹*Augmented Reality* is a special kind of Virtual Reality technique. While Virtual Reality presents the user with a computer only generated environment, Augmented Reality just overlays our real world perception with additional computer generated information. In practice, a surgeon would wear a semi-transparent head mounted display system, which would for instance overlay the surgeon's view of the patients vessel's with colored markers that the radiologist has prepared to indicate hidden calcifications inside the vessels. Please refer to section 4.13 for further information on Augmented Reality

these adapters are optimized to, are not always suitable for visualization needs. As these graphics adapters are built for gaming applications, they are best at rendering a rather small number of polygons with texture maps on them. Visualization often needs a huge number of polygons with simple shading strategies rendered at real time frame rates. This drawback of cheap graphics adapters is vanishing with time, faster systems are capable of rendering higher numbers of triangles at reasonable frame rates. Nevertheless, sometimes graphics adapters provided by different vendors feature less software compatibility as they claim. Even if adapters have been certified to be OpenGL compatible (which is a widely established standard rendering toolkit), hidden differences do exist. Software drivers are implemented differently. Some implementations do not obey the OpenGL specification to every detail. For this reason graphics adapters of different vendors are not easily exchangeable without cumbersome (and hard to find!) modifications to the software.

The problems stated above have been found valid for a lot of hardware upgrade tasks. It seems desirable to change certain configurations of the chosen hardware system, whenever a more powerful option becomes available. For this reason the software system should be designed to be easily adaptable, when it comes to these changes. The best strategy is a strict separation of basic program parts and interface modules for hardware specific tasks. To pick up the example introduced above, it would be a good decision, to pack all OpenGL-specific program calls into separated modules. These modules could be easily modified and exchanged, when a different graphics adapter is to be used.

The hardware system should constantly be upgraded. Although it would avoid a lot of problems to stick to the same hardware configuration and do upgrades just on a not so frequent basis, it is nevertheless necessary. When we decided on the hardware configuration to be used, we picked a graphics adapter developed by a large, well established company. The system has been built around the very powerful features of this system. One year later, the company changed the technology and the board we were relying on was neither produced nor shipped anymore. A lot of effort had to be put into the change of the software system in order to make it easily adaptable to new hardware components. For this reason, the availability of the used hardware system should be of highest concern. If it cannot be guaranteed, the software system has to be built in a flexible way.

When deciding on the hardware platform for a medical workstation, one decision is not avoidable: Should a dedicated graphics workstation or a “general purpose” computer system be used? We have decided to use a standard Intel based PC as the platform for our workstation for the following reasons:

- PCs are cheaper (price, service, maintenance). The total cost of ownership is lowered.
- The evolution cycles of PCs are shorter than for workstations. It is easier to utilize state-of-the-art hardware.
- Hardware add-on devices (like virtual reality hardware, special input devices, special rendering cards (see section 4.9)) are available for PCs.
- As PCs are more established than dedicated workstations, update cycles for hardware drivers, operating system bug fixes, and standard software systems are shorter.

SpaceVision	Intel Pentium III 1GHz	1.5GB RAM
Virtuoso	SGI O2 MIPS R12000 400 MHz	512 MB RAM
Advantage Workstation	UltraSPARC 60 2x450MHz	2GB RAM
Vitrea	Intel Pentium III 2x850MHz	1GB RAM
Plug'n'View	Intel Pentium III 400MHz	512 MB RAM

Table 4.2: Hardware platforms

Nevertheless, high end PC systems have to be used. As medical visualization is usually dealing with large amounts of data, the performance of the system's data buses is very important. The fastest options for hard-disk and memory technologies should be chosen. For "number crunching" tasks, as they often arise in medical visualization (iso-surface extraction, filter operations like morphological operations, etc.) it is desirable to utilize parallelism by employing multiple processors.

Table 4.2 gives a comparison of the different hardware platforms chosen by the developers of the other workstations in this comparative analysis. Siemens has chosen the Silicon Graphics (SGI) O2 platform three years ago, maybe because the easily extendible memory architecture (Unified Memory Architecture, UMA) seemed promising. Today, even SGI is abandoning their family of MIPS-based graphics workstations and is focussing on PC-based systems. GE is using a Sun UltraSPARC platform for the Advantage Workstation. This architecture yields reasonable performance, especially with the dual configuration of processors. Nevertheless, a PC-based system delivering the same performance is significantly cheaper.

The conclusion concerning the hardware platform is: keep the software as open and easily adaptable as possible (Java, clean code, separated hardware specific parts, etc) and exploit all technological improvements as soon as they are available.

4.3 Data Interface

As medical visualization is usually dealing with data acquired from medical scanning modalities, a medical workstation has to be able to interface the patients data. *DICOM*² has been established as a communication standard for such purposes. Various other standards exist. Some modality vendors use proprietary data protocols.

It has to be differentiated, whether a workstation will be connected to a larger information network (such as a PACS system) or whether it will be used as a stand-alone system. We believe, that a medical workstation should be as much integrated as possible. Users usually demand to be able to access patient

²<http://medical.nema.org/dicom.html>

information along with the acquired image slices. It is important for them to compare former studies to the current ones. Especially former diagnosis reports are of interest. For these reasons, it is very convenient, that a possibility to interface information from a PACS system, a radiology information system, or a hospital information system exists while for instance doing a diagnosis on a 3D representation of the patients features.

Another important fact not to be neglected about data interfaces is the bandwidth capability of the interfacing network structures. As the amounts of available data are constantly increasing (see table 4.1), the times for loading the data from database servers which are usually employed for short and long term storage of the patients' data also increase. It is important to keep loading times in a reasonable range in order to preserve the usability of the system in emergency situations or even daily practice. Optimized loading strategies (load on demand, progressive or hierarchical loading, data compression) and high performance interface hardware have to be employed.

All workstations in this comparison feature the possibility to use the DICOM protocol for data access. Additionally, the SpaceVision and Vitrea systems are capable of loading data volume represented by "raw" strings of byte code or a set of single image files in standard formats (like JPG or TIFF). The Virtuoso and the Advantage Workstation are able to interpret proprietary data formats used by Siemens or GE modalities. For SpaceVision, this kind of data has to be converted by means of a dedicated "gateway" software.

4.4 Volume Reconstruction

As medical visualization is dealing with 3D volumes rather than with 2D slices, the first step in the working process is the definition of how slices should be reconstructed into a volume. Although usually this definition is straightforward due to the nature of the scanning process (planar, parallel slices are scanned in sequential order), special demands might arise, which can be dealt with in this initial volume-reconstruction step:

- For certain investigations the set of slices acquired by the modality features a non-uniform inter-slice distance. Examinations of the head are for instance scanned with varying slice distance: In the region where a tumor or aneurysm can be most likely expected slices are scanned with a smaller distance. In order to keep the radiation dose for the patient low, the other parts of the head are scanned in a coarser resolution, i.e., with a larger slice distance.

Such data sets are usually reconstructed into volumes with a regular slice distance. As it is crucial not to lose any resolution and therefore information for the diagnosis in the important parts, usually additional slices are reconstructed by means of interpolation in the low resolution regions of the scan.

- Sometimes it is desirable for the radiologist to remove irrelevant information from the data set in order to reduce the data volume or visually emphasize important features in the volume. A feasible approach is to present the user with a combination of the original 2D slices and a 3D "preview" rendering where he is able to cut down the data set to a *volume*

of interest. Only the specified sub-volume will be used for rendering, all other parts will be omitted.

- For certain applications it is not only desirable to reduce the data set to a cuboid sub-volume, but to remove arbitrary regions of the data set. *Volume sculpting* [WK95] is a technique to incorporate this. The user is for instance able to
 - remove the table of the CT-scanner where the patient was situated on during the examination
 - draw free-form outlines on the slices or even the 3D projection of the data set which prohibits the regions outside the contours from being rendered
 - use sculpting tools in the 3D representation of the data set.

Although this functionalities should also be provided in the rendering features of the software system later-on in the work-flow, it is desirable to get rid of unnecessary parts of the data as early as possible in the work process.

- Another procedure which should be incorporated in the volume reconstruction stage is *volume filtering*. Noise reduction techniques can be employed to increase the quality of the data and get rid of noise artifacts which could spoil the 3D rendering. Thresholding could be used to remove data values which are outside the meaningful scanning levels of the used modality (and therefore can be considered as noise). Other possibilities of filtering include improvement of the visibility of fine structures, which would not be visible very clearly during 3D rendering, like fine structures in the lung or fine ruptures in the inner regions of bones.
- When a set of large-distance slices is to be reconstructed into a volume with a higher resolution in scanning direction, interpolation is used to generate the additional slices. If this interpolation procedure is done in the reconstruction step and not during rendering (on-the-fly) *special interpolation schemes* can be used. If time is a limiting factor only simple techniques like trilinear interpolation can be employed. More computationally expensive approaches like *shape based interpolation* yield smoother and more realistic approximations of the regions missed inbetween the slices during scanning. Figure 4.1 shows some examples for different types of volume-reconstruction approaches.
- Other types of *preprocessing* for rendering can also be accomplished during the volume reconstruction step. Gradient estimations as well as curvature magnitudes can be calculated and stored for rendering. If all preprocessing calculations are performed in a single step at the beginning of the visualization work-flow, it has psychological advantages for the user. The user is not interrupted in his work when he is finally using different rendering and investigation techniques for the diagnosis. Our customers reported, that they would rather like to wait for a longer time in the beginning of the work process than later-on. It does not pay to have the initial idle time reduced by doing the preprocessing "on demand" during the work process.




Figure 4.1: Different approaches to volume reconstruction. Left: linear interpolation; Middle: spline-based interpolation; Right: numerically optimal interpolation.

All workstation in this comparison support the possibility of volume sculpting. Although the approaches are quite different in terms of user interface and usability, basically sculpting by means of contour projection is employed by all workstations. The possibility of advanced volume filtering does only exist in the SpaceVision system. Although the Virtuoso and the Advantage Workstation systems support volume smoothing, no other filtering techniques can be used. SpaceVision additionally supports noise removal, background detection and automatic table removal. SpaceVision is also the only workstation to support the reconstruction of higher order representations of the volume data. Signal and gradient reconstruction can utilize linear, cubic, or spline based interpolation. Additionally, a sophisticated technique for optimal shading can be employed: 4D linear regression [BNKG00].

4.5 Registration and Image Fusion

Another feature of crucial importance in the practical application of medical 3D visualization is the combined visualization of different data sets. Multimodal visualization is dealing with the combined display of data sets of different origin.

It is for instance desirable to merge CT- and MR-data sets in order to compensate for the drawbacks of the two techniques. In CT-data sets bones and structures of high physical density are represented in high contrast whereas soft tissues are hardly detected by the scanning process at all. In MR-data sets on the contrary, soft tissues are resolved to high contrast resolutions, whereas structures that contain little or no water are not visible. Multimodal visualization would extract the advantageous characteristics of both techniques and display a combined representation, which consists of hard structures taken from the CT scan and soft tissues from the MR scan.

Another example for multimodal visualization from the field of neuroradi-




Figure 4.2: Image Fusion Display: 2D overlay of MR (brain) and CT (skull) data (left), and 3D display of the same data sets (right)

ology includes the combination of MR scans of the brain with functional brain scans like functional MRI (fMRI), SPECT (Single Photon Emission CT), or PET (Positron Emission Tomography). Another application of multimodal visualization would visually compare data sets of the same person acquired at different periods of time, which allows judgments on the evolution of tumor structures. Finally, multimodal displays could demonstrate the deviation of a patient's anatomy from a reference model taken from a 3D anatomical atlas.

In order to come up with a combined display, the different data sets have to be spatially aligned. This process is called *registration*. It has to be determined, whether affine transformations (rotation, translation, and scaling) are sufficient to spatially align the data sets. For some applications non-linear transformations like 3D-volume morphing are necessary. MR data for instance features non-linear distortion due to inhomogeneities in the magnetic field of the scanner. Various approaches dealing with these demands have been developed in recent years [Bro92, MV98].

From the user's point of view it is most desirable to utilize a fully automatic registration procedure integrated into the visualization system. Furthermore, the registration process should not waste a lot of the user's expensive time. Finally, the accuracy of the registration should be within a very low error threshold. Unfortunately, there is no technique, which meets all these demands at once.

For the special application of the registration of data sets of the patient's head we found the most reasonable approach to be the usage of external markers. Small spheres made of a material detectable in the different modalities employed are fixed to a mounting frame which is itself fixed to the patient's palate by means of a vacuum system for the acquisition of head scans. These marker spheres can be tracked and registered autonomously by the system resulting in sub-voxel accuracy.

If no markers are available for registration, other feature based [M⁺96] or mutual information [MCV⁺97] based approaches have to be employed. Being

SpaceVision	manual: landmark-point specification; automatic: mutual information and feature based registration
Virtuoso	manual: landmark-point specification
Advantage Workstation	common surface recognition for CT and MR
Vitreia	–
Plug'n'View	–

Table 4.3: Approaches to the registration of multimodal data

computationally very expensive, these calculations should be done during the workstation's idle time, for instance during the night hours.

Finally, there is always the option of having the user mark some characteristic points in all data sets, which can be used for registration. This approach is not only time consuming for the user, but also the accuracy of the registration is limited to the accuracy of the user's point specification.

When data sets are properly aligned, they have to be displayed in a single graphical representation. This process is called *Image Fusion*. Again several approaches have been developed up to now. The visually most attractive methods combine the display of single features of the data sets extracted by means of segmentation. So for instance the shape and location of a brain tumor (MR scan) could be displayed in relation to the vessels in the brain (CT Angiography scan), along with the skull (CT scan). Additionally some functional information on brain activity (fMRI scan) could be color coded into a representation of the brain surface (MR scan). Depending on the available segmentation techniques (see section 4.4) the composition of such a visualization can be quite time consuming.

Less sophisticated approaches simply overlay the different data sets in the different color channels of the graphics system. This technique might be sufficient for diagnostic purposes using just the 2D slices of the data sets. In 3D representations it is quite hard to interpret the information in different color channels.

Other approaches like *screen-door* or *checkerboard* techniques divide the image space in several regions which are used to generate a "woven" overlay image of two images rendered with the two data sets. Figures 4.2 and 4.3 show examples. The user is able to interactively translate the woven structure, in order not to miss smaller details, which are present in one data set but simply covered by a section of the other data set. Although this technique has been proven to be very useful in clinical practice, the number of possible data sets is usually limited to two. If more images are to be combined in this way, the result becomes visually rather disturbing.

Registration and fusion have gained enormous value in clinical practice. Our investigations with customers have shown that for the registration process techniques with minimal time consumption in terms of calculation time and user interaction time have to be employed. For fusion a set of various techniques has to be provided. Depending on the application (diagnosis or communication) the user will want to use different approaches.

SpaceVision, Virtuoso, and Advantage Workstation support the registration of multiple data sets by the manual specification of landmark points common to




Figure 4.3: Image Fusion Display: 2D checkerboard display of MR and CT data, showing three orthonormal slices

SpaceVision	2D and 3D screendoor and checkerboard display, color channel overlays
Virtuoso	checkerboard display in MPR mode
Advantage Workstation	2D checkerboard displays, alpha-blending fusion
Vitrea	–
Plug'n'View	–

Table 4.4: Approaches to the combined display of multimodal data (image fusion)

all data sets (please refer to table 4.3). SpaceVision and Advantage Workstation additionally feature automatic registration. Advantage Workstation is deriving registration information from the alignment of surfaces, which are extracted from the data sets. This approach is limited to the usage of CT and MR data. SpaceVision supports the automatic registration using markers, mutual information, and feature based techniques.

For the fusion of registered data sets, similar approaches are chosen by the systems (see table 4.4). All support the display of checkerboard overlays for 2D images in multi-planar reconstruction (MPR) mode. SpaceVision additionally features the display of 3D checkerboard overlays as well as color channel blending. The Advantage Workstation also features the possibility to merge the data sets into a single data set by means of combining the data values by alpha-blended overlay. The resulting data set can be used like a generic one for further investigation tasks.

Vitrea and Plug'n'View do not support any registration or image fusion at all.

4.6 Segmentation

A crucial part of every medical 3D software system is the possibility to highlight and emphasize single features of the data sets. Using a color classification scheme based on the data values only, it is only possible to separate different data ranges. With CT scans different tissue types can be differentiated depending on the correlation of data values to density properties. If different features within the same data range are to be displayed independently (e.g., brain and the outer skin of a patients head), segmentation tools have to be employed. With MR the situation is even worse: there is no direct correlation of data values to densities. Refer to figure 4.4 for a visualization, which shows different parts of a data set in different colors although all structures are in the same range of data values.

Another very important application for segmentation is the composition of a visualization by showing or hiding certain structures. For a Maximum Intensity Projection display, which is usually used to investigate vascular structures, it is crucial to remove any bones from the data set, which would occlude the view of the vessels.

The final and probably most important diagnostic application for segmentation is the quantitative, volumetric analysis of structures. In the area preceding 3D visualization, the change of pathologic structures like tumors was judged on a relative basis: the tumor increased or decreased over a certain period of time. With the help of very accurate segmentation techniques, an absolute judgement becomes possible, e.g., the size of the tumor decreased by 32 mm^3 in the last two weeks. Other fields of application for quantitative volumetric evaluations include the size of aneurysms, the volume of the lung in case of cancer, and the judgement on calcifications in the patients vascular system.

In recent years a large number of segmentation techniques have been developed in order to deal with these demands. There is no single technique, which is able to deal with all or even the larger part of these applications. We learned, that the tradeoff between generality and accuracy for certain applications can be disastrous. Specialized techniques for the single applications are necessary.

Furthermore, there always has to be a tradeoff between accuracy and the




Figure 4.4: MR data: Although the skin and the outer brain surface are in the same range of data values, they can be rendered in different colors, if segmentation techniques are applied to the brain.

amount of user interaction. Fully automatic techniques work fine for some specialized applications. In other fields, no appropriate automatic techniques exist. The accuracy of the segmentation result highly depends on the way the user is able to assist the system in case of ambiguities. It is most important to optimize these semi-automatic or even manual techniques in terms of their usability. Time is the crucial factor in clinical practice. Depending on the amount of user time required, a 3D visualization system will be used or not.

Table 4.5 gives an overview on the techniques we find most suitable for the different applications. Descriptions of the stated techniques are to be found in section 2.4.

A listing of the different segmentation techniques employed by the workstations in this comparison is given in table 4.6. Plug'n'View and the Virtuoso workstation are lacking sophisticated segmentation techniques. Vitrea and Advantage Workstation are providing specialized segmentation techniques for specific tasks like vessel or colon segmentation. SpaceVision is providing a different approach: the user is given a comprehensive toolbox of segmentation tools, which can be combined to get the best results for any application.

4.7 Reformatations

Scanning modalities acquire the image slices in a certain direction. With CT technology, this direction is always longitudinal to the main patient axis, i.e. head to feet for body regions or along the extension of limbs. Data sets of this type are called *axial* data sets. For MR technology no technical restrictions to the scanning direction exists. Data sets may be scanned in axial, *coronal* (slices are parallel to the patients chest plate), or *sagittal* (sideways) directions.

Application	Technique
<i>Neuroradiology</i> : white and grey brain matter, tumors, functional activations	statistical approaches like Region Growing, semi-automatic techniques like 3D Watershed Segmentation, and manual techniques like Intelligent Scissors or Live Wire segmentation
<i>Vascular Radiology</i> : finding center lines of vessels for Virtual Angioscopy, staging the size and location of calcifications, judging on the volume of aneurysms or stenoses	Vessel Tracking approaches, geometrical analysis techniques
<i>Orthopedics</i> : emphasis or removal of single bones	Threshold based approaches like Seed-Point Segmentation
<i>Pulmology</i> : quantifying the volume of the lung; Virtual Bronchoscopy	statistical techniques like Texture Matching or Region Growing
<i>Mammography</i> : finding microscopic calcifications	statistical texture analysis approaches

Table 4.5: Different segmentation techniques most suitable for different applications

SpaceVision	Region Growing, Vessel Tracking, Thresholding, Automatic Bone Removal, Morphological Operations, Watershed Segmentation, manual tools
Virtuoso	Region Growing, manual tools
Advantage Workstation	automatic edge detection, Vessel Tracking, Morphological Operations
Vitrea	Vessel Tracking, automated colon segmentation, manual tools
Plug'n'View	manual tools

Table 4.6: Segmentation techniques

Once a data set is acquired in a certain direction it is often impossible to repeat the scanning procedure in order to acquire data sets in additional directions. Radiation harm or inconvenience for the patient are limiting the scanning time. For these reasons one of the first 3D applications that was broadly established was the reformatation of data sets. Several approaches to resample the data set in different ways are discussed in the following sections.

4.7.1 Multi-Planar Reconstruction

Multi-Planar Reconstruction (MPR) is dealing with the interactive display of image slices in a new sampling direction from an existing data set. The “classic” employment of MPR would allow the conversion of, e.g., axial data sets into sagittal or coronal ones. The way existing software systems are providing MPR has been found to be insufficient for clinical practice: 2D-resampling planes are defined which display an interpolation of the voxel values in their direct proximity. Using this approach it is for instance virtually impossible to display tiny vascular structures that curl in a real 3D shape. It is crucial to provide a certain “thickness” of the resampling plane, which is able to cover a larger region of this small structures. Using this approach, it obviously has to be defined, how the voxels in the defined neighborhood should be condensed into a single pixel of the resampling plane. Suitable ways to do that are the employment of averaging (resulting in a volume-rendering like representation) or choosing the maximum or minimum (resulting in a MIP or MINIP (minimum intensity projection) like display).

Additionally to simply changing the scanning direction of the data set into one of the two other main directions aligned to the axis of the data volume (this is called *Classic MPR*), it should of course be possible to define arbitrarily angled scanning directions. In order to structure the way the users work with MPR it should be clearly differentiated between *Oblique MPR* (the resampling direction may only be rotated around one axis) and *Double-Oblique MPR* (the resampling direction may be rotated around two axes, i.e., the axis is pointing arbitrarily through the data set).

4.7.2 Curved-Planar Reconstruction

Whereas MPR is dealing with the resampling of the data set on planes, Curved-Planar Reconstruction allows resampling according to curved structures. Figure 4.5 illustrates the concept. As most anatomical structures are curved in one or the other way (like blood vessels or bones), ways to reconstruct images of such structures are desirable in a 3D-medical workstation. The most common applications of such techniques include the reconstruction of a curved image, which displays the (curved) human spine in it’s entirety, or the simulation of dental panorama x-raying (a single image showing all the teeth in relation to the jaws). Figure 4.6 shows another example featuring a curved blood vessel.

Due to the visual complexity of the reconstructed views, there has to be a straightforward way of working with CPR. We found it to be most intuitive for the user to simply click into any of three orthonormal MPR views, defining a path which follows the curved structure which is to be reconstructed. For the sake of the visual continuity of the reconstructed image (and for the sake of the




Figure 4.5: A curved planar reconstruction (CPR) allows the resampling of volume data following curved structures of interest [Kan01]



Figure 4.6: A curved planar reconstruction (CPR) is used to display a cross section of the peripheral artery of the right leg in its entirety [Kan01].

users convenience), the path points should be used as control points for a spline-interpolation scheme, which smoothes the path. While the user is defining path points, two additional views should be displayed. Immediate feedback to the accuracy of the location of the defined path points is provided by a real-time display of the flattened CPR surface. The user should be given the possibility to correct the location of single points anytime by just dragging the points with the mouse. The changes to the CPR view should of course be displayed in real-time while the user is dragging the points.

For blood vessels alternatively segmentation techniques can be employed, which are capable of finding the centerline automatically (refer to section 2.4.5).

4.7.3 Reformatting the data set

The MPR approach described above allows the user to move sampling planes through the data set in order to view reconstructed image slices. As radiologists are used to work with image stacks it is also desirable to give them the possibility to calculate entire sets of new images from the original data set. This process is called *data-set reformatation*.

The simplest application of reformatation would calculate a set of parallel slices in an arbitrary sampling direction. The parameters for the resampling process should be user-definable to the same extent as at the acquiring modality, i.e., the user has to be in control of parameters like slice-thickness, slice distance, the way the region of slice slabs are integrated into 2D slices, and finally the number of slices.

In addition to parallel slices other ways of arranging the slices should be provided: orthonormal to a path (calculating true axial slices of the spine or a blood vessel), extruding from the path (covering a vessel), rotational (lung examinations), rotational around a path (vessel-stenoses display).

It is very important to provide the possibility to store the reformatted images in the PACS or the employed data-base system. Radiologists need to employ all the tools they are usually using to read generic images: changing the windowing level, zooming, employing filters, etc.

All the workstations in this comparison provide the functionality of Multi-Planar Reconstruction (axis-aligned as well as double-oblique modes). Comparing the speed of the recalculation of slices, all workstations provide interactive feedback. The Plug'n'View system and the SpaceVision station achieve the highest frame rates, followed by the Advantage Workstation and the Vitrea system. The Virtuoso system seems to provide the least performance, nevertheless, feedback is still interactive. Additionally, all systems provide the possibility of Curved Planar Reconstructions. The usability of the specification scheme for the path used by the CPR differs widely. The Advantage Workstation features the approach that can be used most effortlessly. The Virtuoso uses a specification scheme, which takes some training to be used efficiently. The approaches of the other workstations work well also for untrained users. The SpaceVision workstation also provides the possibility to automatically generate animation sequences of CPRs rotating around the defining path.

All workstations provide the possibility of reformatting the data set in different ways. The speed of calculation is comparable for all systems. Only the Plug'n'View system is lacking the possibility to export the reformatted data sets in DICOM, which would allow further usage in a PACS network.



Figure 4.7: Measuring the grade of a stenosis in the contrast agent enhanced blood vessel (at the bottom of the image).

4.7.4 Measurements

With the rise of MPR and reformations the need for good measurement tools has evolved. It is finally possible to measure the maximal diameter of a blood vessel by moving a MPR plane to get the largest possible cross-section and employ a simple distance-measurement metaphor. Accurate numerical calculations in high-precision numbers are necessary in this context.

Obviously three different dimensions are necessary: length, area, and volume. While user interaction for length and area measurements is pretty straightforward, the problem of volumetric measurements needs to employ means from the field of segmentation. Please refer to section 4.6 for details on volumetric measurements by means of segmentation tools.

Let us consider the problem of measuring the diameter of a vessel stenosis in order to demonstrate quite a difficult problem in medical visualization, which also largely influences the topic of measurements. A radiologist would do (using a medical 3D workstation) the following: he would use a MPR display to find a representative cross section of the stenosis (see figure 4.7 for an example). Then he would use the mouse to draw a line from one side of the wall of the vessel to the other side. Due to the partial volume effect there is no clear transition of the vessel lumen. Depending on the radiologist's judgement of the vessels extends, the measurement might vary by 25%.

The computer system might help in this case. We found it useful to integrate the following mechanism: The user is specifying two points in the vessel, one in the critical section of the stenosis and another one in a non-pathological region of the vessel just close to the stenosis. The system automatically finds axial cross sections and also the vascular slices encircling the user specified points. As in the two calculated cross sections the transition zone usually is comparable, the relation of the size of the two vessel cross sections can be calculated rather accurately. The relation is the important fact for the radiologist. Most reports read like this: The degree of the stenosis located 6 cm below the knee is 25,4%.




Figure 4.8: Simple surface based visualization (left) versus volume rendering (right hand side). Volume Rendering is able to utilize transparencies, colors, and segmentation masks.

With this, quantitative comparisons are possible with the help of a sophisticated 3D workstation.

All workstation in this comparison feature the basic functionalities for measuring lengths and angles. Only the Advantage Workstation, the Vitrea system, and the SpaceVision station provide tools for the automatic or semi-automatic analysis of pathological vascular structures like stenoses.

4.8 Rendering

As medical visualization is concerned with the generation of meaningful images from a stack of image slices, the most basic part of every medical 3D workstation is the rendering subsystem. A variety of different approaches to the display of 3D-volume data have been developed over the years [Den90, Rob93]. Two major branches have evolved: *Volume Rendering* (concerned with the display of the entire volume using color and transparency classifications) and *Surface-based Rendering* (using polygon based surface geometries extracted from the data set for rendering). See figure 4.8 for a visual comparison.

Although these approaches are different, they have one thing in common: There always is a tradeoff between speed and quality. These issues will be discussed in detail in the following sections.

4.9 Direct Volume Rendering

Technologies providing very fast rendering frame rates (like VolumePro or Shear Warp Factorization) are usually limited in terms of rendering resolution and interpolation quality. High-quality rendering approaches (like ray tracing) are too slow to be used in clinical practice. Sometimes progressive approaches are used to overcome these problems. Starting with a low image resolution (low enough

to be rendered with real-time frame-rates) better resolutions of the visualization are rendered as soon as the user stops manipulating the display, e.g., stops rotating the displayed object.

Much time is used here on the generation of many intermediate quality levels, which mainly fulfill the purpose of keeping the user interested while rendering is still in progress. We chose a different, two-step only process. Using the highest quality, real-time rendering technology available today (i.e., the VolumePro hardware-rendering accelerator [PHK⁺99]), we provide the user with a possibility to view and manipulate the volume-data set in feasible quality. If higher quality is needed, the user would select a viewpoint using the VolumePro display and start the calculation of the high quality rendering in software by clicking a button. All calculations are done transparently for the user, using additional CPUs (in case of a multiprocessor computer) or other CPUs in the PACS network running idle. Therefore the user is able to go on investigating while the high-quality renderings are calculated. Via small iconic representations of the requested images the user is able to see the progress of the rendering process. When rendering is finished, the user is able to bring the high-quality representations to the screen and have a closer look. It is very important in this stage to provide the possibility to “synchronize” the view point of the interactive VolumePro display to the viewpoint used to calculate the high-quality image. Users usually want to see that little feature they just found in the image from a little different viewing angle.

Rendering speed is just one factor to the usability of volume rendering. Different ways of employing this powerful tool for clinical practice have to be included in a medical workstation.

The Virtuoso and the SpaceVision workstation are employing the VolumePro technology for interactive volume rendering. SpaceVision is additionally capable of rendering high-quality high-resolution visualizations by ray casting techniques (not interactively). The Advantage Workstation and the Vitrea system are using a software-based rendering technique, which progressively renders visualizations of improved resolution and quality. The Plug’n’View workstation is also using a software-based rendering technique, which produces very clear images at high frame-rates. Unfortunately, higher resolutions cannot be achieved.

4.9.1 3D Windowing

When Volume Rendering is employed for the display of a volume-data set, a mapping of data values to optical properties like color and transparency has to be defined. The simplest definition uses a ramp function for mapping data values to transparency values (refer to figure 4.9). As real time rendering feedback is provided, the user can be given interactive control over the location and the width of this ramp. The visual appearance of the changing rendering inspired one of our collaborating radiologists to the term “3D window level change”. Although incorrect from a theoretical point of view, the analogy to the well established mechanism of changing the “window level” of a 2D generic image is valid.

By changing the width and location of the transparency ramp in the range of data values the user can browse through the different tissues of the data set. Basic diagnostic procedures are already possible in this rendering mode. If for the sake of an easier communication with referring colleagues different types




Figure 4.9: Using a linear ramp function to map data values to transparencies. No color information is used.

of tissue have to be visually distinguishable, more elaborate transfer functions have to be employed.

4.9.2 Using transfer functions

In the case of CT scans a transfer function defines the mapping of transparencies and colors to different types of tissue. For MR scans different regions in the range of signal intensities are mapped to colors and transparencies. Visually appealing renderings can be created with the help of transfer functions (see figure 4.10). Nevertheless, the specification of a transfer function raises several challenges. An optimal transfer function has to be chosen from a potentially infinite number of possible ones. This search space is obviously too large to be searched in detail by the user.

We found the optimal specification scheme to be a combination of suggestions calculated by the system and a certain amount of user interaction for the fine-tuning of the proposals. The employment of the VolumePro technology enables us to give real-time feedback on the user's modifications to the transfer function. Chapter 7 describes this approach in further detail.

Additionally, it is important to give the user the possibility to store transfer functions, he found useful for his purposes. Especially with CT data it is often possible to adapt a stored transfer function to a new clinical situation with just minor modifications.

All workstations in this comparison are capable of using transfer functions for volume rendering. Transfer functions can be stored and reused for other data



Figure 4.10: A transfer function has been used to map the data ranges for skin, bone, and the contrast agent in the blood vessels to different colors and transparencies for this CT scan of a human head.

sets. The Advantage Workstation and the Vitrea system are using a gallery approach to display the results of using stored transfer functions with the investigated data set. The Vitrea system is using a very simple approach for the specification, where only a color transition and a single opacity ramp are used for the transfer function. Although this yields satisfying results for most applications, the control of the user over the transfer function is very limited. The Virtuoso workstation is using a very sophisticated user interface for the specification, which requires quite some training by the user, but does not feature any advantages at all. The SpaceVision workstation is able to use the specification scheme described in chapter 7 additionally to the standard specification schemes used by the other workstations.

4.9.3 Using indexed volumes

When the results of a segmentation process are to be visualized (see section 4.6), different transfer functions have to be assigned to the regions covered by the different segmentation masks. This approach is called *indexed-volume rendering*. Figure 4.11 shows an example. Currently, indexed volumes are not supported by the current generation of VolumePro technology (VP500). Nevertheless, it is no problem to include segmentation results in the high quality renderer of our software system (see section 4.12).

For the interactive rendering method a different strategy had to be included. Surface representations of the segmentation masks are rendered together with the volumetric display in order to give the user a visual impression on the location of segmented features. Until the next generation of VolumePro technology is released, this hybrid-rendering work-around is employed.




Figure 4.11: Indexed-Volume Rendering: The intervertebral disks and the spinal chord have been segmented and are displayed together with the vertebrae.

Only the Advantage Workstation is capable of displaying indexed volumes at interactive frame rates in this comparison.

4.9.4 Tools

Together with the volume-rendering display a set of tools has to be provided in order to enable the user to prepare the visualization to the needs of diagnostic or communicational purposes.

First of all it is important to integrate the possibility of *volume sculpting*[WK95]. The user often needs to remove parts of the volume data set which obscure the view to relevant parts of the patient's anatomy. The most straightforward approach is the application of *cutting planes*. Using an interactive display it is desirable to have cutting planes which remove all data values on one side of the plane which can be positioned by moving the mouse. Such cutting planes should be provided in different ways of angular orientation: either aligned to the main axis of the data set (axial, sagittal, or coronal) or aligned to arbitrary cutting directions. Special cutting planes can be provided, which selectively remove data values depending on the existence of segmentation masks. Figure 4.2 shows an example of a cutting plane, which was used to remove the top part of the skull but does not affect the brain.

Additionally to the planar removal of parts of the data set, a tool for the removal of arbitrary cutting regions should be included. An intuitive way of applying such a tool is the following: The user defines a cutting contour either by defining a set of cut points, which resemble the cutting contour when connected, or the user just uses the mouse to outline the desired region with an arbitrary continuous contour. When the cutting action is triggered by the

user, the contour is extruded in viewing direction in order to define the cutting volume. Eventually, the user has to decide, whether the data values inside or the data values outside of the cutting volume are to be removed, i.e., either the cutting volume should be removed or everything but the cutting volume should be removed. Using this interaction paradigm the complex definition of a 3D-cutting volume is simplified to the definition of a 2D-cutting contour. The user is able to remove arbitrary parts of the data set by repetitive rotation and contour cutting.

Another very important tool for the composition of the visualization is a light source, which can be moved interactively by the user. Certain delicate structures in the patient's anatomy, e.g., fine fractures, are not visible under direct lighting conditions. Tiny dents or fine fractures become only visible, if a lighting direction non-orthogonal to the surface normal vector is used. For this reason it is crucial to provide the user with the possibility to interactively manipulate the lighting direction when looking for such structures.

All workstations in this comparison provide the possibility of removing parts of the visualization by means of axis-aligned cutting planes or arbitrarily defined cutting regions.

4.9.5 Maximum Intensity Projection

An important type of volume rendering is *Maximum Intensity Projection* (MIP). Used for the diagnosis of vascular structures, it has been one of the first well established techniques of 3D visualization in the medical field. Up to now, the calculation of MIP images was rather time consuming, which prevented the interactive application of this tool. Animation sequences were calculated off-line, which gave the user an impression of the 3D structure of the vessels, but obviously did not allow the manipulation of the view point after the calculation of the sequence in case an important detail was not clearly visible in the animation. With the development of optimized software techniques and dedicated rendering hardware, MIP has become interactive. The user has become able to rotate, zoom, and pan the display with real-time frame-rates when looking for structures of diagnostic relevance.

Additionally different variations of the "classic" MIP have been developed. Due to the nature of the rendering technique there is no depth information present in a MIP image. For this reason, animation (either pre-calculated or interactive) is necessary to get a 3D impression of the rendered structures. The 3D shape cannot be derived from still images. The introduction of *depth shading* is able to soften this problem. Structures farther away from the view point are displayed in darker color, thus introducing the third dimension even to still images. Another approach to overcome this problem is called *Local MIP*. With this approach the maximum along a viewing ray is displayed only for the local maximum that is encountered first, which results in darker borders of curved objects. These darker borders indicate the overlapping configuration in the viewing direction thus introducing 3D information into the visualization. Figure 2.4 gives a comparison of the three different techniques. Local MIP and Depth Shaded MIP are not supported by rendering hardware like VolumePro, but software approaches capable of interactive frame-rates exist and should be integrated into every module of a medical 3D workstation dedicated to the diagnosis of vascular structures.

All the workstations in this comparison feature the interactive display of Maximum Intensity Projections. SpaceVision and the Virtuoso workstation are using the VolumePro hardware-rendering technology, whereas the other workstations are computing the visualizations by means of software calculation. Again, only the SpaceVision workstation is capable of rendering visualizations of arbitrarily high resolutions. Depth-shaded MIP and Local MIP are only included in the SpaceVision workstation. For these rendering techniques that are not supported by VolumePro software calculation is used, which yields a slightly worse quality/speed trade-off.

4.10 Surface rendering

The display of surfaces extracted from the volume-data set is well established since the dawn of medical visualization. *Surface Shaded Display (SSD)* is the common term for this approach. By defining a threshold the user specifies where a surface should be extracted from the data set. Based on the threshold definition methods like *Marching Cubes* or *Marching Triangle* extract an iso-surface from the data set. As the surface extraction procedure is computationally expensive, optimization techniques like *Propagational Marching Cubes* [KKI99] should be employed. The triangular mesh of the surface can be rendered by means of graphics libraries like *OpenGL*. In order to achieve feasible rendering times, mesh decimation algorithms can be employed to reduce the number of triangles in the surface. For the interactive exploration of mesh representations of the patients anatomy, *levels-of-details* [SS97] approaches can be employed.

The main advantage of SSD over volume-rendering techniques is the inherent pre-segmentation of the representation due to the surface extraction. Connectivity information can be used to identify single objects in the displayed geometry. In our software system we provide the user with the possibility to select single objects by simply clicking on them. Selected objects can be deleted or assigned a distinct color. Using this paradigm the user is able to quickly compose the visualization to his needs. Sometimes it is desirable to manually break the connection between two objects in order to assign different colors or in order to delete just one of them. An approach similar to the contour cutting technique described in section 4.9.4 can be used to achieve this goal.

The main disadvantage of SSD is the problematic dependency on the threshold definition. An iso-surface is extracted wherever a transition of two tissue types in the volume equals the threshold value. As there are partial volume effects and smooth transition present in all volume-data sets, the size and volume of the extracted geometries highly depends on the definition of a feasible threshold. Please refer to figure 4.12 for a comparison concerning the influence of the iso-value on the extracted surface. This is especially a problem with anatomical structures, which have been enhanced by means of contrast agents. If the contrast agent is not distributed homogeneously, the structure cannot be extracted by means of a single threshold. Interval-based approaches have been developed to overcome this problem.

Finally, it is desirable to enable the user to display multiple iso-surfaces (extracted for different thresholds) simultaneously. Using this approach the spatial relation of structures of different types of tissue can be displayed.

SpaceVision is the only workstation in this comparison, which is utilizing




Figure 4.12: Comparison of two iso-surfaces extracted from a CT angiography data set at two different thresholds: an iso-value of 1180 was used for the surface on the left hand side, an iso-value of 1140 for the surface shown at the right hand side

polygonal meshes for the display of surfaces. All other workstations are generating surface visualizations by means of volume rendering using transfer functions with a single discrete opacity transition. SpaceVision is capable of exploiting the advantages of the representation of structures by meshes. Single objects can be picked for manipulation by utilizing connectivity criteria. Furthermore, structures can be exported in various 3D geometry formats in order to be used by prototyping-software systems for prosthesis planning.

4.11 Hybrid Rendering

The previous section described the suitability of surface rendering for the display of structures which are represented by well defined transitions in the data volume, e.g. bone structures or vascular structures enhanced by contrast agents. Volume Rendering is superior to Surface-based Rendering when structures of not well defined size (like tumors) are to be visualized. *Hybrid Rendering* combines the advantages of both technologies. Surfaces can be rendered along with parts of the volume-data set classified by a transfer function.

As there is no hardware support for Hybrid Rendering up to now, computationally expensive software implementations are usually not capable of achieving interactive frame rates. The Advantage Workstation is using such an approach to produce visualizations nearly at interactive frame rates. For the SpaceVision station Hybrid Rendering is only applicable for the creation of high quality renderings described in the next section. All other workstations in this comparison are not providing the possibility of Hybrid Rendering.

4.12 High quality rendering

As the quality of interactive visualizations is usually limited in terms of image resolution and signal reconstruction effort, it is desirable to integrate into the medical workstation the possibility to create high quality renderings. Computationally expensive rendering techniques like *Ray tracing* can be employed, if the generation time of the visualizations is of minor concern. Additional features aiding the visual power of the images can be integrated:

- *Mirrors* may enhance the information present in the image, when placed in strategically optimal positions. Using this approach for instance not only the front side of an anatomical structure can be shown, but also parts facing away from the viewing direction [AHH⁺94].
- *Magic Mirrors* even enhance the values of classical mirrors. Magic mirrors are not restricted to the physical limitations for the transport of light. Displays originating from different rendering techniques can be used for Magic Mirrors. So for instance, a different transfer function or a contour display could be used instead of the reflection of an object (refer to figure 5.6).
- Another feature suitable for the integration in a high quality rendering module is the usage of sophisticated transfer function definitions. Not only the data values should be used for a color classification, but also higher-order derivatives of the acquired data signal, like *gradient* and *curvature*. When the gradient is taken into account, homogeneous regions in the data set, which should not contribute to a direct volume rendering display can be omitted, thus emphasizing transitions between different types of tissue. A curvature-based transfer function can be used to enrich the display with color coded information on the local curvature of structures, in order to visually mark structures of a certain shape. With this approach for instance polyps (sphere-like shape) and lesions (ellipsoid shapes) can be emphasized in the display of a patients colon. Transfer functions based on higher-order reconstructions are computationally too expensive to be used in a hardware rendering system, therefore this feature is realized in a software based high-quality rendering-module.

The application of transfer functions can be spatially limited by separating planes. For applications targeted to communication it is sometimes valuable to use different transfer functions in spatially disjoint regions of the data set. An example for this approach would be the composition of a scene, where the upper part of a patient's head has been removed entirely but for the brain by means of a cutting plane (figure 4.2). Different colors could be used for the brain surface and the patients skin in the lower part, if the rendering system is aware of the two separating-plane dependent transfer functions.

- For the application of reconstructional surgery it is valuable to allow the introduction of artificial objects into the visualization. The rendering system should allow the simultaneous display of the patients anatomy and a three-dimensional model of the prosthesis which is going to be applied. When planning the surgical procedure, it is possible to align the prosthesis

with the patients anatomy and to check the size and shape of it even before the real prosthesis is crafted. Additionally, the surgeon's tools can be displayed together with the patients data in order to check the suitability of different tools in the different stages of the surgical procedure.

As there is still no hardware-supported hybrid rendering available, special strategies have to be employed in order to provide the user with the possibility to move and manipulate tools and prostheses with respect to the patient's structures in real time [vWHVP91, Lev90].

- When preparing still images or movies in order to present especially interesting cases or configurations of the patient's anatomy, *annotations* are usually used to explain the cases. Usually a textual annotation together with an arrow indicating the spot of interest is used. Nevertheless it is visually way more appealing to place a 3D model of the text and the arrow at the corresponding location of the patient's model and have it rendered altogether. The same approaches utilized for the integration of tools and prostheses mentioned above can be employed here.
- For the preparation of high quality presentations it is always necessary to integrate the possibility of placing measurement tools to the patient's structures. A display of a tumor for instance is meaningless without a visual hint on its size. One possibility to indicate the scale of structures is to place a 3D model of a classical ruler next to it. More accurate displays can be achieved by wrapping the structure into a mesh with a rectangular grid on it. As the size of rectangles is preserved, the size of the displayed structure can be easily derived by counting the rectangles.
- For presentations there should also be the possibility to combine the 3D display of anatomical structures with multi-planar reconstruction (MPR) planes or curved planar reconstruction (CPR) surfaces. If parts of the 3D model are hiding the reconstructions, transparency can be used to increase the practical value of the composition.
- Sometimes cutting planes are not the tools of choice when it comes to the necessity of removing parts of the visualization. When for instance a blood vessel shall be partially extracted from the data set, a cutting plane is not able to adopt the curved shape of the vessel. For such cases the possibility of *Volume Sculpting* [WK95] by tools of arbitrary shape should be provided. The user would remove unwanted parts of the data set by moving a sphere- or cuboid-like tool in the display to chisel away structures interactively. An enhancement of this functionality should be transfer function or segmentation-dependent sculpting-tools, where the user would specify beforehand, which parts of the data (in terms of data range or segmentation results) should be affected by the sculpting procedure.

As the main purpose of the generation of high quality renderings is the task of communication, it is even more desirable to produce movies and animation sequences. As the techniques and additional functionalities described above are not capable of rendering the visualization in real time, a two step procedure should be employed:

1. The user plans the animation sequence. Using interactive rendering techniques, a low-quality preview of the final animation sequence can be checked versus the demands of the user.
2. The system renders the animation sequence in a batch processing mode.

For the planning of the animation sequence, a feasible metaphor for the specification of the varying parameters has to be included in the medical workstation. The most common parameter to change over time is the location and orientation of the camera. Options for the specification of camera paths include:

- Simple rotational movement can be easily specified. The user moves the camera in the interactive preview mode into the starting position of the animation. In order to define the rotational movement of the visualized scene, only three parameters have to be specified: the axis of rotation (with respect to the camera coordinate system), the angle of rotation, and the number of frames, which should be rendered within this angle. The rotation angle does not necessarily have to be specified in terms of the number of degrees, but preferably by simply rotating the interactive preview display to the desired end position of the rotation. Using this approach, it is convenient for the user to specify, whether the rotation of the camera should be a rotation around the position of the camera or rather a rotation around the viewed object with the viewing direction fixed to the center point of the object.
- A translational movement is also specified very easily. Again, the user moves the camera into a starting position. Then it is assumed, that the movement is restricted to the viewing direction. The user specifies the end point of the translation by moving the camera to this position.
- By allowing the combination of a number of rotational and translational movements, any path for the animation can be defined.
- Another possibility for the specification of the animation path is the usage of a spline-based path-definition. The user specifies the control points of the spline curve and checks the resulting camera flight in the preview display.
- As the developed 3D workstation features a scripting interface, the most powerful way to define the path of the animation is to code the movement and rotation of the camera manually. As this approach requires a basic understanding of programming and of 3D coordinate systems, it is not suitable for all users.
- For special applications like Virtual Endoscopy, the path for the animation can be found automatically by the system. See section 2.5.1 for details.

Other parameters that change over the time of the animation can be specified using similar approaches: transfer function definitions, position and orientation of cutting planes, display parameters for segmentation results, etc.

As the rendering of the final animation sequences is very time consuming, different approaches can be employed to speed up the process:

- Multi-processor computers can be used for the number-crunching task of rendering [Wes95, PC92, Zen97].
- As in a hospital's PACS system there are usually more computers than users at a given time, a concept originating in Hollywood's movie industry can be employed: the task of rendering the animation is shared by several machines in the hospitals network. All computers running idle can donate CPU cycles for the calculation of animation frames. In such a *render farm* there has to be a central authority, which checks for idle machines and distributes the work accordingly. For this task a set of network daemons has been developed. Of course, this approach is only applicable in the multi-computer scenario of a hospital. In the case of a single stand-alone 3D workstation, it can not be applied.
- If long animation sequences of high quality are to be rendered, it is always desirable to run the process in a batch mode during the night hours. If results are needed immediately, the render farm approach should be used or a dedicated render server should be employed in order to not stall the 3D workstation for the process of rendering the animation sequence. The 3D workstation is only used to define the parameters of the animation, other computers are used to render it.

The SpaceVision workstation is the only one in this comparison to provide the tools and functionalities described above. All of the tools are realized in the scope of the high-quality rendering-module. The interactive rendering display utilizing the VolumePro hardware-rendering is only used for preparational tasks. The Vitrea system and the Advantage Workstation provide a possibility to render animation sequences, too. Nevertheless, they lack the possibility of generating visualizations and animation sequences of high resolutions.

4.13 Virtual Reality devices

With the introduction of virtual reality techniques new interaction tools for the investigation of many types of data became available. Immersive environments can enable a user to deal with 3D representations using more natural interface paradigms. Of special interest to the field of medical visualization are *Augmented Reality* systems [SFGS96]. This technology does not immerse the user in an entirely artificial, computer generated environment, but rather overlays the user's real world view with complimentary information using see-through head-mounted displays. So for instance, a surgeon is using an augmented reality system to overlay his view of the patient's structures with visual aids and guiding information, which he has prepared based on radiological data prior to the surgical procedure.

For this application, the usage of a *Head-Mounted-Display* system is desirable. An augmented view is realized by means of two semitransparent LCD-displays which the user wears mounted in front of his eyes. In order to overlay the view of the real patient's structures with reconstructed 3D representations, a *tracking system* is necessary to align real and virtual data. In many virtual-reality systems, magnetic tracking systems are used to accomplish this task. Nevertheless, under surgery conditions this technology cannot be used due to

the influence of the magnetic fields on other surgical apparatuses. Optical tracking systems utilize multiple cameras and markers on the tracked user's head in order to overcome this drawback [Sei99].

Even in radiology departments it is often desirable to provide the user with an enhanced 3D view. Stereo viewing systems like shutter glasses or 3D monitors provide a better 3D impression of the visualized structures. Although these viewing aids are not likely to be used by radiologists (who are trained to "see" 3D structures by mentally reconstructing visual information), these stereo viewing systems are mainly used for communication purposes. Surgeons, clinicians, and patients prefer interactive 3D displays to still images.

All of the workstations in this comparison are capable of utilizing tracking devices for 3D input. All but the Plug'n'View workstation are capable of rendering stereo views to be used by shutter-glasses displays or virtual environments. None of the workstations features a module for the intra-operative overlay of reality and computer generated visualizations. Such modules are under development for the SpaceVision workstation, the Virtuoso system, and the Advantage Workstation.

4.14 Storing the results

An integrated 3D workstation has to provide options for the storage of resulting reconstructions. Specialized storage strategies have to be integrated for the different types of results:

Images – The storage of still images is straightforward. If the workstation is integrated in a PACS network, still images can be appended to the patients' examinations data. Storing the images in the PACS archive is the easiest way to provide the results to referring doctors, as any viewing workstation in the PACS systems is capable of displaying the still images.

Additionally, it is desirable to enable the user to store images also on his local hard disk. Different graphics file formats should be provided, because depending on the application (sending via E-Mail, presentation on a web page, including in a lecture notes or slides for a talk, etc), different levels of quality and resolution are necessary. We experienced, that most users feel inconvenient when being forced to use special graphics software systems for the simple task of converting image data to another resolution or format.

Movies – Animation sequences can be handled just like still images. Storing movies as a sequence of single images to a PACS archive enables referring doctors to view the movies in standard viewing workstations by animating the sequence in a so-called "cine loop". For storing movies to the local hard disk, different movie file formats should be available. So if for instance a very long movie is to be written to a CD-ROM, often a reduction in quality is desirable.

3D models – For two reasons it is also desirable to enable the user to store extracted geometries (e.g., iso-surfaces or sparse volume data sets) in the PACS system or to the local hard disk. Firstly, if the work on one patient has to be interrupted, it can be continued at a later point in time without

SpaceVision	2D snapshots via DICOM, JPG, TIFF, BMP; movies via DICOM, MPG, AVI; 3D models via Quick-timeVR, DXF, 3DS
Virtuoso	2D snapshots via DICOM, JPG, BMP
Advantage Workstation	2D snapshots via DICOM, JPG, TIFF, BMP; movies via single images
Vitrea	2D snapshots via DICOM, JPG, TIFF, BMP; movies via AVI
Plug'n'View	2D snapshots via JPG and TIFF

Table 4.7: Different possibilities for storing visualization results

having to repeat all necessary preparational steps. Secondly, it is possible to provide a referring doctor or the patient with a small stand-alone viewing software, which allows the interactive investigation of the preprocessed geometry even on standard computers. Obviously the possibility to explore the data by interactively manipulating the camera position and other rendering parameters is superior to just watching movies or looking at still images.

A special aspect in the scope of storing the results is the fact, that referring doctors are not necessarily located in the same hospital as the radiologist. With the rise of the internet, fast and reliable data transmission has become possible. Nevertheless, two possible drawbacks have to be overcome:

- The bandwidth of the transmitting networks is limited. Especially with large volume data-sets sophisticated strategies have to be employed for internet-based visualization. Various research projects have tried to tackle these problems [MLG98].
- Security hazards are involved with the transmission of sensitive patient data. Strategies like anonymization or encryption have to be used.

For the sake of reproducibility and other legal issues, a 3D workstation should be able to save the entire "session", which represents all the actions a user took to generate visualizations. This includes the current state of his work along with the history of all changes, like the choice of certain parameters of the visualization. By saving this information, it is possible to continue working right a the point of progress, when work has been interrupted. Furthermore, it is possible from a legal point of view to investigate how and why a radiologist came up with an incorrect finding.

Table 4.7 gives a comparison of the different strategies employed by the workstations in this comparison for storing resulting visualizations.

4.15 Conclusions

In this chapter the most important aspects have been discussed, which need to be taken in consideration, when developing a medical 3D workstation. First of all, it is most important to keep the system independent in terms of hardware

and software. The research field of visualization is a vivid one, permanently producing new techniques and technologies. When it comes to the integration of new features it is invaluable to have an open system. Standard data access schemes should be integrated. When the workstation is part of a PACS network, DICOM is the communication protocol of choice. With the introduction of high-resolution imaging techniques, the reconstruction of missing information in volume data sets has become less important. Nevertheless, sophisticated techniques for the reconstruction of gradient information and higher order derivatives like curvature have to be integrated. As all imaging techniques have advantages and pitfalls, a possibility of volume registration and fusion is highly desirable. The combined display of CT, MR, and functional information can provide crucial information to the diagnosis of special diseases.

When pathological structures are to be highlighted or MR-data sets are the basis for the visualization, segmentation techniques have to be employed. As no universal technique to satisfy the diversified demands is available, a variety of specialized automatic, semi-automatic, and manual segmentation techniques should be integrated. Well established tools like Multi-Planar Reconstruction, Curved-Planar Reconstruction, and arbitrary reformations are considered to be basic functionalities of every medical 3D workstation. Advanced reconstruction and rendering techniques like Maximum Intensity Projection, Surface Shaded Display, and Volume Rendering have to be provided for different applications. Real-time-rendering frame-rates are necessary to give immediate visual feedback to the user and provide an intuitive working environment. High quality or high resolution renderings can be produced off-line without the necessity of user interaction or supervision. Special fields of application like Virtual Endoscopy, planning of surgical procedures, the training of students, and the presentation of radiological findings are the first areas in medicine, where 3D visualization has been employed.

The most important requirements when building a medical workstation include usability (providing an easy to use environment, which helps to prevent human errors), reliability (ensuring that no artifacts have been introduced and that no crucial details have been omitted), and reproducibility (using deterministic approaches to guarantee the comparability of different results).

Five different workstations have been compared in this chapter. It is hard to come up with an overall conclusion concerning the value of the different systems. The Advantage Workstations seems to be the most evolved system. The reason for that might be that it is the workstation in this comparison with the longest history of development. Although the SpaceVision workstation features more functions and tools than the Advantage Workstation, it still needs some polishing in terms of usability and workflow. The Vitrea system provides a "quick-start" approach, which makes it easy to come up with impressive visualizations with a few mouse clicks even for unexperienced users. Unfortunately, the systems does not provide any advanced possibilities for controlling the outcome of the visualization procedure. The Advantage Workstation and the SpaceVision system are providing functions for the different levels of experience of users. Beginners as well as advanced users will find tools and approaches to meet their needs. This is not true for the Virtuoso system. It requires quite some experience in the field of visualization to come up with satisfying results. Although this system provides some powerful tools, a rework in terms of usability seems to be necessary. It is hard to compare the Plug'n'View system to the

	Pro	Con
SpaceVision	has the most comprehensive set of tools and functions; very fast hardware rendering	needs improvement of workflow
Virtuoso	very fast hardware rendering	low usability
Advantage Workstation	most evolved system in terms of work-flow integration and technical quality of functions	would need faster rendering
Vitrea	easy to use for unexperienced users	does not satisfy "advanced" users
Plug'n'View	affordable	is lacking a lot of crucial functions

Table 4.8: Conclusions concerning the workstations in this comparison

Rank	Usability	Speed	Functions	Quality	Prize	Overall
1.	AW	SV	SV	AW	P'n'V	SV
2.	Va	Vo	AW	SV	Va	AW
3.	P'n'V	P'n'V	Vo	Vo	SV	P'n'V
4.	SV	Va	Va	P'n'V	Vo	Va
5.	Vo	AW	P'n'V	Va	AW	Vo

Table 4.9: (Very subjective) ranking. AW=Advantage Workstation, SV=SpaceVision, P'n'V=Plug'n'View, Va=Vitrea, Vo=Virtuoso. The overall ranking (rightmost column) has been found by averaging the other properties.

others, because it is a very new workstation, which is still under development and features a lower number of functions and tools than the other workstations. Nevertheless, these functions are among the best there are in terms of speed and usability. If the development is continued in this kind of quality, the Plug'n'View system will be among the best of the market in a few years. Table 4.8 gives an overview on the conclusions concerning the different workstations and table 4.9 gives a ranking.

4.16 The future of 3D medical visualization

Visualization will change the way assistants and radiologists work. More and more tools will be developed, which either help to diagnose and treat health problems better and faster, or enable doctors to investigate diseases and injuries, that could not be diagnosed with conventional 2D approaches. As the amount of information available on the patients' anatomy is drastically increasing with the further development of image acquisition technologies, *Computer Aided Diagnosis* will play an important role in the future. Radiologists will not have to investigate all the available images and data anymore. The system will

guide them to regions, where their experience has to be employed. Additionally further specialization of the different radiological disciplines will occur. With internet-enabled diagnostic tools, remote diagnosis by specialists will be daily practice.

Chapter 5

Multiple Views and Magic Mirrors

This chapter is based on the following publications:

- Helmut Doleisch. **Multimodal Visualization of Anatomical and Functional Data of the Human Brain.** *Diploma thesis*, Institute of Computer Graphics and Algorithms, Vienna University of Technology, Austria, 1998.
- Andreas H. König, Helmut Doleisch, and Eduard Gröller. **Multiple Views and Magic Mirrors – fMRI Visualization of the Human Brain.** *Spring Conference on Computer Graphics (SCCG) 1998*, Budmerice, Slovakia.

5.1 Introduction

Over the last decade, the use of digital modalities for clinical diagnosis and therapy planning has dramatically increased. New developments in three-dimensional volume acquisition techniques such as Spiral CT (*Computed Tomography*) scanning, faster and better pulse sequences for MRI (*Magnetic Resonance Imaging*) and improved functional imaging provided by modalities such as PET (*Positron Emission Tomography*), SPECT (*Single Photon Emission Computed Tomograph*), EEG (*Electro EncephaloGraphy*), MEG (*Magnetic EncephaloGraphy*), and fMRI (*functional Magnetic Resonance Imaging*) are creating a rapidly increasing demand for integrated multimodality 3D visualization. It aims at efficiently presenting information from multiple modalities, in order to obtain a comprehensive picture of clinical cases. Therefore, depending on the clinical requirements, it might often not be sufficient to consider vascular, anatomical and functional information separately but to superimpose images from different modalities.

Functional MR (magnetic resonance) imaging is used to determine regions of the human brain, which are responding to sensory stimulation, motor or cognitive activity or emotional states. Neuroimaging experts investigate such scans usually by employing their talent to visually extract the characteristics of activated regions from the scanned volumes. This exploration task is usually done




Figure 5.1: Conventional approaches to multimodal visualization [RLF⁺98]. Top row: selected MRI slices are merged with functional data by pixel interleaving [RSA⁺94]. Middle row: the maximum intensities of the data are projected onto axial, coronal and sagittal planes bearing brain outlines and millimetric Talairach coordinates [TT88]. Bottom row: the maximum data intensities are projected along the line of sight onto the cortical surface.

on twodimensional slices of the human brain, where the functionally activated regions are highlighted in some special way (see Fig. 5.1, top row). Viewing several slices at once (side by side) is referred to as the *image galleries* approach. This way of investigating the data sets provides detailed 2D anatomical context information. Nevertheless, this way of dealing with multiple data sets does not aid the user in forming a three-dimensional model of the underlying spacial relationship of functional activations and the anatomical geometry. Other approaches try to increase the amount of presented information by utilizing *maximum intensity projections* (MIP), mapped onto 2D reference grids with brain outlines [Fri94] (see figure 5.1, middle and bottom row). These representations aid the task of side by side comparisons, but nevertheless provide just very rough information of the anatomical context. Therefore it is very hard for the user to mentally reconstruct the three-dimensional structure of the patients brain.

Three-dimensional surface renderings of the brain, where the functional activation is encoded by coloring the cerebral cortex provides a misleading anatomical context. Activated subcortical foci below the surface of the brain are obscured. Furthermore, all foci seem to lie on the surface of the cortex, as no depth cues are provided.

Direct volume rendering would be able to overcome this problem. Assigning transparencies to data ranges would allow the deeper structures to be visualized. Therefore direct volume rendering would be the most usable approach to the solution of this problem. Usually, however, anatomical detail has to be sacrificed, when employing transparencies [HP96], decreasing the visual appearance of details significantly.

For these reasons K. Rehm and other authors recently proposed, that the anatomy of the brain should not be visualized three-dimensionally at all

[RLF⁺98]. They use symbolic representations known from other fields of visualization (like fluid or information visualization) to encode the location and size of foci. In order to place the foci into anatomical context, they are projected onto the walls of a cube like environment together with different representations of the brain. We are going to utilize a similar approach of adding information to the visualization in Section 5.4.

Direct volume rendering still seems to be the most useful approach to clearly give visible hints on the location and size of activated regions. We are going to present some extensions to overcome the problem of reduced visibility in the following sections.

First of all, the given problem of our special case of combined visualization of functional and anatomical data of the human brain will be stated in detail. The problems arising when using standard direct volume visualization techniques will be discussed in section 5.3. The first improvement using multiple views will be discussed in section 5.4. Additional aids for the investigating user (Magic Mirrors) are introduced in section 5.5. A results section will provide images and statistics to prove the usability of our techniques. Finally conclusions will be drawn.

5.2 The problems encountered in multimodal volume visualization

In this chapter a volume visualization technique for multimodal data of the human brain is presented, consisting of functional data of the human brain (measured by fMRI) and anatomical data of the brain (measured by different methods of MRI). Other types of multimodality are out of the scope of this paper and are not discussed. The functional data represents the activation potential of the brain at a certain point in time.

The test situation used was the following: A patient has to perform the repetitive action of thumb and index finger tapping. The activation potential of the brain is measured. This gives the functional context of the reaction to the action of the fingers. The fMRI scans are acquired during stimulation as well as under resting conditions. By calculating the differences of these images the activated regions of the brain can be obtained. For a complete description of MRI specialized literature [CS92, SNN95] should be consulted.

The problems encountered with multimodal volume visualization of anatomical/functional data of the human brain include the following:

- *the data sets have different resolutions*: the resolution of the functional data sets is usually significantly lower than that provided by the anatomical data. This is due to the fact, that the functional data has to be measured interactively. An anatomical MRI scan is acquired using more time.
- *the spatial extents of the data sets are different*: This is again, due to the time constraints in fMRI measuring. Usually only parts of the brain (for example only 5-7 slices) are scanned for the functional data set.

Therefore the data sets first have to be spatially aligned through *registration* [MV98] (see chapter 4.5). Afterwards visualization is used to identify the

regions of the brain, where the changes in the activation potential occurred. This visualization step is the main point of this chapter.

5.3 Index volumes for multimodal visualization

As mentioned above, the usage of shaded surface displays does not serve our purpose of visualizing multimodal data very well. First of all, inner structures of the brain (like white matter) are obscured by the covering surface. When activated regions are encoded by coloring the according parts of the surface, incorrect information is presented: activations seem to appear only on the cortical surface. The parts within the brain responding to activation are not visible at all. For these reasons we are not able to use any surface based visualization techniques like for instance *isosurface extraction by marching cubes* [LC87].

Direct Volume Rendering (or probabilistic rendering) is concerned with the visualization of inner structures. Data values are mapped into color and opacity values by the usage of transfer functions. During image rendering physical properties of light transportation in non-opaque media are simulated. This approach serves our problem of multimodal volume visualization very well. Different methods in direct volume rendering are available:

- *Splatting* [Wes90] is an object space approach, which composes the image by accumulating the contribution of every single voxel in the data set.
- *Ray casting* [Lev88] is an image space approach, which accumulates contributions from voxels encountered along viewing rays. As this method approximates techniques used in photorealistic rendering (global illumination), the simulation of physically correct phenomena like refraction or reflection is possible. Having a special enhancement to our technique in mind (described in Section 5.4), ray casting is the technique of our choice.

For the chosen approach of ray casting, a special scheme for the specification of transfer functions had to be employed. *Transfer-function volumes* are used to specify different transfer functions for separate regions of the data set to be visualized. A transfer-function volume is a three-dimensional array storing an index to a set of transfer functions for every single voxel. This index is used to access a transfer function for mapping the data value of the voxel into the physical properties color and opacity. When the number of transfer functions to be used is limited, e.g. to 4 or 8, the indices can be packed more densely and significant amounts of storage space are saved.

Assigning transfer functions on a voxel per voxel basis has several advantages:

- A special transfer function is defined, which should be used to encode the activated regions of the anatomy. During the *registration* process, the index of this transfer function is assigned to all voxels which are at least partially intersected by a voxel of the functional data set with activation potential. During rendering of the anatomical data set activated regions are colored using this special transfer function and are highlighted.
- Using a "NULL" transfer function (mapping every data value to zero opacity) *volume sculpting* facilities can be realized quite easily. Setting regions




Figure 5.2: Different problems arise when using direct volume rendering. When surface rendering is employed, the regions of activation are hidden by the cerebral cortex (left hand side). Right, very low opacity allows deeper structures to be visualized. However, it is hard to tell, where the foci reside exactly.

of the transfer function volume to the index of this special transfer function, various "cut-away" possibilities can be realized, because voxels whose values are represented by "NULL" transfer functions simply do not appear in the visualization. These regions can be specified by cutting planes, block cut operators, or even on a voxel by voxel basis. The anatomical data set used in the scope of this work has been segmented in this way. A segmentation mask gained by manual segmentation of the anatomical data set specified the voxels of the non-brain regions of the head scan, which were rendered using a "NULL" transfer function. The original data set has not been modified in any way.

- When this approach has to be optimized for the calculation time of the ray-casting step, the data volume can be converted into a color/opacity volume using the transfer-function volume in a preprocessing step. If rapid changes of the used transfer functions are desired (for instance when rendering several frames for an animated visualization), the transfer-function volume does not have to be changed at all. Simply the set of transfer functions referenced by the indices is modified.

Direct volume rendering is capable to overcome the problem of obscuring. Nevertheless, significant detail in image quality is sacrificed when employing low opacity contributions of the voxels. Figure 5.2 shows a comparison of the two extremal cases. When a transfer function is used, which maps the voxel values of the anatomical data set into high opacity values, the visualization yields a surface-like representation (figure 5.2, left). When very low opacity values are used, hardly any anatomical context is visible (figure 5.2, right). It is not easy to describe the location of activated regions according to the visualizations in figure 5.2. A solution to overcome this problem is shown in the next section. To compensate for the loss of image details, additional information is brought




Figure 5.3: The principle of placing image planes as multiple views of the scene for enhancing volume visualization: without (left) and with (right) textures added.

into the visualization by enabling alternative views of the brain.

5.4 Multiple Views - Enhancing the visual impression of direct volume renderings

From a single still image it is very hard to get the correct visual impression of depth. In the presented technique the image produced with the transfer-function volume is enhanced by adding three different views of the object to the final image. These multiple views are realized by positioning three image planes behind and below the object. Images displayed on this image planes are again generated by using the ray-casting technique [Lev88]. These three image planes are added to the scene as shown in figure 5.3. The white parts on the three planes (see figure 5.3, left) depict the location, where the additional images are displayed as textures (figure 5.3, right). Each of these images has its own set of camera parameters. In the example of figure 5.3 the view points for the image planes are behind (below, respectively) the planes. These image planes provide multiple views and are quite different in their properties from, e.g., mirrors or shadow walls. Placing a real mirror at the same location (an approach used by A. Kaufman [AHH⁺94]) would result in completely different images. Nevertheless, this different behaviour was aimed at for the further developments discussed in the next section.

The benefit gained from these multiple views is manifold:

- The location of activated regions becomes obvious even in still images, as the three-dimensional context of the two data sets can be recognized more easily when multiple views on the same scene are available. This is especially the case for highly transparent renderings (compare the image on the right hand side of figure 5.2 with figure 5.4, left hand side).
- The problem of obscuring is solved rather satisfactorily. In Figure 5.4, right hand side, for example a block of the anatomical data has been




Figure 5.4: Left: Using multiple views, the exact localization of activations is also possible with transfer functions using transparency; Right: Even occluded regions like the cutaway part become obvious with the help of multiple views

cut out at the backside, but due to the multiple views provided by this technique, the locations of activations are still visible.

5.5 Magic Mirrors - Illustrating additional information

The benefits of multiple views of the brain generated by volume rendering can even be enhanced. As the employed image planes of the multiple views do not show the real-world physical behaviour of, e.g., reflecting planes, several quite different rendering techniques can be used to generate the textures. Leaving the domain of physical reflection completely, the extended version of the multiple views is called *Magic Mirrors*.

- *"Projection"* vs. *"Reflection"*: The multiple views textures used in the last section were created by ray casting the brain anatomy with a view point located behind (below, respectively) the image planes. For this reason, the quasi-reflective property of the views was gained. It is easy to locate the view point for texture generation on the other side of the brain geometry. Utilizing this approach, projections of the brain in direction of the mirror can be generated. The main views onto the brain (front, top, right) can be used to enhance the visual impression of the threedimensional representation of the brain (figure 5.5, bottom mirror).
- Different transfer functions can be used for each of the Magic Mirrors. Various characteristics of the multimodal scenario can be emphasized in a single image. Figure 5.5 demonstrates the relation of activated regions to white and gray brain tissue.
- Additional views to the anatomy created by techniques other than direct-volume rendering can be used to provide complementary information. In




Figure 5.5: Magic mirrors employing different transfer functions and "Projection"

figure 5.6 a 3D "contour map" of the brain and activated regions (generated by the ray-tracing approach) was mapped onto the Magic Mirror at the right hand side of the scene. The usage of other techniques like *maximum intensity projections*, *orthonormal cuts*, *color encoded cuts*, or *segmentation maps* as Magic Mirrors is also possible.




Figure 5.6: Magic mirrors with contour display

5.6 Results

The goal of the work presented in this chapter was to develop a technique, which allows the identification of the location of regions with activation potential. In order to prove the usefulness of the presented Magic Mirrors approach a non-clinical test scenario was set up.

The data sets used for this work were gained from a 3-Tesla MRI scanner. The anatomical data sets were acquired using a pulse sequence called *MDEFT* (Modified Driven Equilibrium Fourier Transform). These images are T1-weighted. The functional data sets used in this work are T2*-weighted, and have already been pre-processed, so that the regions of activations are emphasized in the volume, and all other regions are neglected.

The data sets featured the resolutions shown in Table 5.1. A small group of test persons was asked to interpret the given Magic Mirrors images and verbosely describe the location as well as the size of activated regions in relation to the brain anatomy. All evaluators reported, that it was an easy task to derive the location of activated regions from the multiple views and Magic Mirrors representations.

The volume of the activated regions was not estimated as easily. Although this information is implicitly known to the system from the properties of the functional data set, it might be desirable to get an instant impression of the size of the activated volume.

5.7 Conclusions

The problem of obscuring when utilizing direct volume rendering in the scope of multimodal volume visualization of functional/anatomical data of the human brain has been dealt with.

Transfer-function volumes are used to realize highlighting and various volume sculpting tasks by assigning different transfer functions on a voxel by voxel basis. Multiple views of the visualized brain geometry in a single image provide additional information. Different representations of the multimodal scenario can be displayed by Magic Mirrors. Various combinations of data representations (utilizing different transfer functions and contour displays) provide the possibility to intuitively derive information like location and size of activated regions from the visualization.

Nevertheless, there were some open problems, which have been overcome with the integration of latest technology. A faster ray-casting system (VolumePro-based, see section 4.9) is utilized, which allows the interactive investigation and manipulation of the visualization. Regions of activation are treated as entities, allowing to visually encode the magnitude of volume size and other properties in our direct volume rendering approach.

	anatomical data set	functional data set
resolution	256 x 256 x 64 slices	64 x 64 x 5 slices
voxel size	1mm x 1mm x 2mm	4mm x 4mm x 4mm

Table 5.1: Resolutions of the data sets used to get the results shown in this chapter

Chapter 6

Visualization of large data sets

This chapter is based on the following publication:

- Andreas H. König, Helmut Doleisch, Andreas Kottar, Brigitte Kriszt, and Eduard Gröller. **ALVis – An Aluminium-Foam Visualization and Investigation Tool**. *Joint Eurographics - IEEE TCVG Symposium on Visualization (VisSym00)*, Amsterdam, Netherlands.

In 1997 the introduction of the multi-slice scanning technology was announced. Until the first full-scale data sets were available (late 1998), we had to find a source of volume data, which would let us allow to prepare ourselves for the problems, which would arise with the huge medical data sets. In material science, industrial CT scanners are used to investigate the internal structure of workpieces. As there is no radiation harm, a high scanning resolution can be used, which results in large data sets that are of approximately the same size as the expected multi-slice data sets. As solid pieces of metal are hard to compare to medical structures, we decided to use scans of a very novel field in material science: metal foams. These semi-hollow structures were comparable to the anatomical cavities of human patients to some degree. This chapter describes the strategies, which had to be developed in order to be able to deal with large-scale data sets.

6.1 Introduction

In recent years there has been a continuous increase in interest for metallic foams. Foams based on aluminium provide a number of advantages over solid materials: low specific weight, high specific rigidity, high energy-absorbing potential, reduced conductivity of sound and heat, non-combustibility, and the ability to be recycled. Metallic foams can be compacted at a relatively constant stress level which qualifies them for energy absorption applications (e.g., car manufacturing). Their mechanical behaviour is influenced by imperfections caused by the foaming process. Spatial variations in the foam cell-size distri-




Figure 6.1: Comparison of visualization (left) and reality (right)

bution¹ and imperfect cell wall geometries are studied as examples for such imperfections. Small samples (about $2 \times 2 \times 2$ cm) are deformed in a repetitive way. Pressure is applied to the sample in one direction, until cells start deforming and the sample loses some small amount of length in this direction. Then the deformation characteristics of the sample is investigated. This procedure is repeated until all the cells are entirely collapsed and the sample is completely compacted. In order to aid these investigations, a tool is needed, which allows the non-destructive investigation of samples during the single steps of the deformation procedure. The next section discusses the demands material scientist expect from such a tool.

6.2 Requirements for the investigation of aluminium foam

The demands for the investigation of aluminium foam risen by material-science experts are manifold. Due to the repetitive nature of the investigation of deformation results, the foam sample must not be destroyed for the investigation. Thus, cutting techniques cannot be employed. The approach presented in this chapter utilizes a non-immersive way of gathering information about the local properties of the foam sample: The sample is scanned with the usage of a computer tomograph (CT) modality. Whereas medical CT modalities are limited in terms of scanning resolution (due to the radiation harm for the patient), industrial CT scanners feature resolutions of about 0.04mm.

The visualization of the scanned volumetric data sets shows the requirements of classical volume visualization problems as well as special demands:

- **Interactivity:** As the investigation procedure includes zooming in on different cells of interest, an interactive visualization approach is crucial.

The manipulation of the view point (movement of the camera) as well as

¹The field of computer graphics often refers to the term *cell* as a sub-cuboid of a volume data set with eight data values at its corners. In the field of metal foam research the term cell refers to the air filled bubbles within the solid material. In the work presented in this chapter the term cell is used in the later meaning.

manipulations of the data set (cutting planes, etc.) should be possible with interactive frame rates.

- **Accuracy:** As the relation of cell size to wall thickness is of interest to the investigation, high accuracy has to be ensured. The visualization should represent the structure and topology of the real sample as accurately as possible (with respect to the limitations of the scanning procedure).
- **Measurements:** Tools for investigating the size of cell features have to be provided.
- **Shape criteria:** Different metrics for the description of cell characteristics are needed.
- **Statistical evaluations:** Besides the investigation of local cell properties, also information about global foam sample properties and the entirety of cells is desirable. Visual aids for the user to distinguish between regions of the foam sample with different characteristics should be provided.

Existing commercial as well as academic visualization systems cannot be employed for the investigation task. Due to the high spatial resolution of the CT scanners, data sets of enormous size are generated. Existing systems are not capable of rendering these data sets in real-time. Furthermore the display of special properties of foam cells as required by the cooperating material scientists (like shape and size criteria) is not supported by existing systems.

The approaches employed in AlVis to deal with these requirements are described in the following sections.

6.3 Interactive Rendering

In order to meet the requirements of interactive rendering, a surface shaded display approach is used. Modern hardware graphics accelerators are capable of rendering large amounts of polygons in real-time frame-rates. The OpenGL [NDW93] graphics library is employed to generate the visualization.

The segmentation of the volume data set into a surface representation is done with the aid of an isosurface extraction approach. The well known *Marching Cubes* [LC87] method is utilized. To overcome the drawbacks of ambiguity artifacts [Mat94, NH91] an approach by Montani [MSS94b] is used.

Three problems have to be overcome with the usage of isosurface extraction: A suitable threshold for the surface generation has to be defined. In order to allow the display of foam cell properties, individual cells are identified. Then a way of dealing with the huge number of generated triangles has to be found.

6.3.1 Threshold definition

The aluminium/air boundary in the sample is represented by a highly wrinkled convoluted surface due to the high frequency nature of the data. Isosurfaces as used in medical visualization applications are usually smoother. For the presented application the specification of an appropriate threshold is crucial in order to meet the accuracy requirements of the investigation procedure. The threshold has to be chosen in a way such that the volume of the foam sample is

preserved. Errors in threshold definition have a cubical influence to the volume of foam cells. A rough approximation of a cell is a sphere. A small change in threshold changes the radius of the sphere. This change has a cubical influence on the volume of the sphere. The appropriate threshold for the extraction of the aluminium/air boundary is dependent on the type of modality and external influences. For these reasons, the utilized threshold has to be adaptively derived for every sample to be investigated.

The following procedure yields the optimal threshold:

- The real sample is weighted using high precision scales. Let W_s be the weight of the sample.
- The specific weight of aluminium is known: W_{Al} is 2.7kg/dm^3 (2.7g/cm^3)
- The volume of aluminium in the sample resolves to $V_{Al} = \frac{W_s}{W_{Al}}$
- The volume of the sample cuboid can be calculated using the measured dimensions of the sample: $V_s = x_s * y_s * z_s$
- The ratio R of aluminium to air in the sample resolves to:

$$R = \frac{V_{Al}}{V_s}$$

- Given the histogram of the volume data set $H = \{h_0, \dots, h_{\max}\}$, $h_x = \{v | \text{value}(v) = x\}$, R_t is calculated as

$$R_t = \frac{\sum_{i=t}^{\max} h_i * i}{\sum_{i=0}^{\max} h_i * i}$$

R_t describes the ratio of aluminium to air in case t is selected as threshold. Selecting the threshold t so that $R_t \leq R \leq R_{t-1}$ guarantees that the apparent volume in the computed sample closely corresponds to the real world sample. This is very important as the lacunarity (size and distribution of foam cells) is crucial for judging structural characteristics.

Visual comparisons and spatial relation measurements on the real sample as well as in the geometry representation of the visualized sample have affirmed the validity of the derived threshold (see Figure 6.1).

6.3.2 Foam cell identification

Before the isosurface is extracted, single cells are identified. The volume data set is interpreted as a binary image, separated into voxels having lower or higher values than the threshold derived in the last section. A *two-pass labeling* approach yields the number of distinct foam cells as well as the different sets of voxels contributing to the single cells. Based on this information, the isosurface is extracted on a cell by cell basis. The generated triangles are stored indexed with accordance to the cell they represent. Special care is taken about additional closing triangles at the cut up faces of the foam sample.

6.3.3 Rendering acceleration

The extraction of the isosurface from the volume data set yields an enormous number of triangles. Even for up-to-date hardware-rendering accelerators, special strategies have to be employed in order to provide real-time frame rates.

Two cases have to be distinguished for rendering: The viewpoint might be inside the boundaries of the foam sample, investigating inner cell structures. It might also be outside the sample, when the user wants to gain an overview of cell size and shape distribution.

Most of the time during the examination the viewpoint will be inside the cave-like structure of the foam. As the maximal viewing distance is very restricted, only a small part of the geometry data is visible. A special subdivision scheme is applied for rendering. Similar to the proposals of other authors [GHY98, WVV94] a regular partitioning of the display geometry is employed. The volume data set is subdivided into sub-cuboids of user defined size. These subregions are called *macrocubes*. The surface portions generated during the marching cubes procedure are stored separately for each macrocube in an optimized OpenGL triangle-strip data-structure. The rendering procedure employs a progressive approach:

- **Interactive phase:** As long as the user moves the viewpoint, the rendering budget is limited by the performance of the used graphics hardware. A limited set of macrocubes within the viewing frustum is determined. This set is rendered front to back, until no more time for this frame is available (see figure 6.2, left hand side). Thus the maximal amount of information for real-time frame-rates is rendered.
- **Completion phase:** When the user stops moving the camera, more rendering time can be invested. The OpenGL z-buffer is checked for regions, where no triangles have been rendered possibly due to not yet rendered macrocubes. Using the camera location and direction, the according set of macrocubes is derived, which is rendered front to back, either until the z-buffer is completely covered with triangles or until no more macrocubes inside the viewing frustum are left to render. The later case might emerge, if the user has positioned the camera facing from within an open cell to the outside of the sample.

Un-filled z-buffer regions can only emerge in situations, when the camera faces structures of connected foam cells, which tunnel into the depth of the visualization. This happens rarely during the investigation. Usually just the interactive phase has to be performed.

When the camera is outside the limits of the data set (to gain an overview of the entire foam sample), the macrocubes rendered in the interactive step are those that lie in the three faces of the data set facing the camera.

In addition to the gain of rendering speed, the usage of macrocubes has an additional benefit. When the triangle sets of different macrocubes are colored individually, the size of foam cells can be quickly visually derived by the user. Due to the usage of perspective rendering the perception of depth also gains a lot by the different size of near and far macrocubes (see figure 6.2, right hand side).




Figure 6.2: Left: The rendering subdivision scheme using macrocubes. In the interactive phase, macrocubes 1-4 are rendered. In the completion phase, macrocubes 5-8 are rendered in order to patch an un-filled region. Right: Individually colored macrocubes allow a quick perception of feature size and depth information. The size of one macrocube in this image is $0.32 \times 0.32 \times 0.32\text{mm}$.

6.3.4 Cutting planes

The user often encounters the problem of obscuring when investigating the foam structure. This problem can be reduced by the usage of cutting planes. It is desirable to have cuts aligned perpendicular to the viewing direction in order to avoid misinterpretations by the user due to perspective distortion. The cutting plane feature is implemented in an efficient way utilizing the standard OpenGL viewing frustum culling. A “near plane” is specified, which prohibits parts of the triangle geometry from being rendered, which are closer to the viewpoint than the near plane. The user defines the distance viewpoint - near plane using a user-interface control metaphor. The removal effect of the cutting plane to the graphical foam representation is displayed in real time, giving the user the possibility to explore occluded parts of the foam sample by moving the cutting plane (see figure 6.3).




Figure 6.3: Left hand side: Sample with cut-away part. Right hand side: Distance measuring metaphor. The current length of the rubber band is 0.6mm.

6.4 Measurements

The size of foam cells is one of the most important properties for the investigating material scientists. It is important to know, whether single larger cells are surrounded by a lot of smaller ones. The size of cells in relation to the thickness of the surrounding walls is characteristic for the deformation behaviour.

In order to allow spatial measurement a metaphor based on the analogy of a rubber band (elastic) has been integrated. A three-dimensional crosshair is used to specify the endpoints of the elastic. For navigational reasons the axis of the crosshair are aligned to the three main axis of the foam sample. In order to specify a point for the measurement, the user moves the camera to the location of interest, until the crosshair penetrates the rendered triangles in this region. Pressing a button fixes one end of the rubber band to this location. Then the crosshair is moved to the endpoint of the measurement, while the rubber band is elongated. The length of the rubber band is calculated based on the voxel size and is displayed in real time during this procedure. Figure 6.3 shows the usage of this tool for the determination of the diameter of a very small foam cell. Using another button, the rubber band can be fixed to the measured points, which allows to check the proper location of the rubber band from another viewpoint.

6.5 Focused Views

Obscuring problems, which cannot be dealt with using cutting planes, have to be solved with special techniques. In order to extract the important visual information, an approach well established in the field of information visualization is employed: *Focused Views*. Not the entire foam sample is rendered, but only cells which meet certain criteria. Additionally, colors can be used to illustrate criteria characteristics of certain cells. Figure 6.4 demonstrates the usage of a simple focused view, where only cells in the inner part of the sample are rendered. This allows the material scientist to gain an overview on the size and shape distribution of cells, which determines the inner stiffness of the aluminium foam. Useful criteria to restrict the rendered information include the size and the shape of single cells. These criteria are discussed in detail in the following sections.

6.5.1 Cell size

The evaluation of cell size as a focus criterion is straight forward. The number of voxels enclosed by the cell is a feasible approximation to the volume of a cell. The user is free to specify an interval of cell volumes to be rendered. Figure 6.5 shows the application for small and large cells. It is convenient to allow the specification of the restriction in absolute (e.g. “smaller than 7mm^3 ”) or relative numbers (e.g. “larger than 80% of the cells”).

6.5.2 Cell shape

One of the requirements of the investigation of the foam sample was the determination of the shape of the foam cells. The uniformity of cells guarantees an optimal stress absorption potential. A metric for the judgement of cell topologies is defined. The desirable shape of a foam cell is a sphere. A shape criterion




Figure 6.4: Application of a cutting plane (left hand side). Focusing (right hand side) renders only cells, which are not cut open by the faces of the sample. Colors indicate the shape factor of cells.




Figure 6.5: Focused view based on size of cells (left: largest 10% of all foam cells, right: smallest 20%, middle: both)

f based on the comparison of the cell shape to a sphere can be defined as follows: Let B be the set of voxels, which are enclosed by the surface of a certain cell C . $B = \{v_i | v_i \text{ inside cell } C, i = 0 \dots n - 1\}$ Let \mathbf{M} be the center of mass of B :

$$\mathbf{M} = \frac{1}{n} \sum_{i=0}^{n-1} \mathbf{location}(v_i)$$

A distance function evaluates the absolute distance of certain voxels to the center of mass:

$$\text{distance}(v_i) = |\mathbf{location}(v_i) - \mathbf{M}|$$

The average distance D_{avg} evaluates to:

$$D_{\text{avg}} = \frac{1}{n} \sum_{i=0}^{n-1} \text{distance}(v_i)$$

Two different shape characteristics are of interest to material scientists investigating aluminium foam:

- foam cells closely resembling spheres, only having a small number of high frequency outlying peaks
- foam cells deviating from the shape of a sphere by a lot of minor dents and bumps

Figure 6.6 gives a graphical representation of these two cell types. Material scientists want to investigate, if one of these common cell types is responsible for high cell collapsibility during the deformation process. Now two different




Figure 6.6: Cells closely resembling spheres although having single peak-like deformations (left hand side) feature high f_{\max} and low f_{avg} values. Cells with a lot of minor dents and bumps (right hand side) vice-versa.

criteria for a cell can be calculated: the average and the maximal distance of voxels to the center of mass:

$$f_{\max} = \max_{i=0..n-1} |\text{distance}(v_i) - D_{\text{avg}}|$$

$$f_{\text{avg}} = \frac{1}{n} \sum_{i=0}^{n-1} |\text{distance}(v_i) - D_{\text{avg}}|$$

It is convenient to have these absolute f_{\max} and f_{avg} criteria scaled to relative values on the basis of all cells in the foam sample. Furthermore the two values can be combined into a weighted sum for the ease of user interaction. The specification of the weights w_{\max} and w_{avg} decides which kind of shape anomaly shall be emphasized or ignored when a focused view is applied.

$$f = w_{\max} \frac{f_{\max}}{m_{\max}} + w_{\text{avg}} \frac{f_{\text{avg}}}{m_{\text{avg}}}$$

where m_{\max} (respectively m_{avg}) is the maximal value of f_{\max} (respectively f_{avg}) for all cells of the sample. Using this approach it is for instance possible for the user to have all cells with a shape criterion f in the best percentile displayed by focusing the view to cells with $f < 0.10$. If cells deviating with single peaks from the shape of a sphere are of interest, w_{\max} should be assigned a high value, whereas if w_{avg} is emphasized, cells with a lot of minor deviations are favored.

6.6 Results

The presented approach was implemented using Microsoft Visual C++ (Developer Studio ver. 6.0), MS Windows NT 4.0. AIVis was tested on an Intel PIII based system (450 MHz, 512 MB RAM, 3DLabs Oxygen GVX1 graphics




Figure 6.7: Focused view based on shape of cells. Blue cells: close to the shape of a sphere, small value of f ; green: medium value of f ; red: large value of f .

accelerator). The data set used to generate the results presented in this chapter featured the characteristics shown in table 6.1. High resolution images of the results presented in this chapter are to be found on the WWW:

<http://www.cg.tuwien.ac.at/research/vis/Miscellaneous/alufoam>.

The presented system has been developed in cooperation with a group of material scientists. User studies involving four aluminium foam experts and three different data sets (each with a number of deformed versions) were made. Permanent feedback during the development has assured that the functionality of the system met the requirements of the experts. Interactivity has shown to be the most crucial point for daily work. Only the usage of the presented focused views techniques enabled the quick perception of local and global homogeneity properties of the foam structure. As a first research result ALVis has already enabled the material scientists to distinguish different classes of foam cells, which might help to optimize the deformation behaviour of the foam.

size of sample	$22.3 \times 22.0 \times 30.0\text{mm}$ (14.718 cm^3)	relative weight of sample	0.499g/cm^3
weight of sample	7,43 g	fill ratio	18.5%
scanner resolution	$0.04 \times 0.04 \times 0.04\text{mm}$	value of aluminium	220
data value resolution	8 bit	optimal threshold	35
data set resolution	$626 \times 626 \times 760$	# generated triangles	9112628
size of data set	298 MByte	# triangles per voxel	0.04
scanning time	5 hours 37 minutes	# ambiguity cases	11077 (0.04%)
		macrocube size	$8 \times 8 \times 8$ voxel
		# identified cells	1150

Table 6.1: The statistics at the left hand side of the table were known before the investigation. Values at the right hand side were determined by the system.

6.7 Conclusions

In close cooperation with material scientists a system for the investigation of aluminium foam has been developed. Between iterated steps of deformation CT scans of a foam sample are acquired. Due to the highly complex and wrinkled structure of the aluminium/air boundary, the visualization of these data sets is difficult. As high accuracy is necessary, the threshold value for the employed iso-surfacing technique has to be volume preserving. It is derived using the weight and size of the real sample. Due to the structure of the foam, enormous amounts of triangles have to be rendered. A sub-division scheme utilizing macrocubes guarantees real-time frame-rates. In order to aid the investigation task, statistical properties of individual foam cells can be displayed. Focused views enable the user to restrict rendering to cells with certain size and shape criteria. User studies with material-science experts proved the usefulness of the integrated tools. As a first research result, certain types of foam cells have been identified, which might influence the deformation behaviour of the aluminium foam.

The most important conclusion is of course, that we have found a way of handling volume-data sets of huge dimensions. When the announced multi-slice CT scanners will start to produce large medical data sets, we are prepared.

Chapter 7

Advanced Transfer Function Specification

This chapter is based on the following publications:

- Andreas H. König and Eduard Gröller. **Mastering Transfer Function Specification by using VolumePro Technology.** *Spring Conference on Computer Graphics 2001*, Budmerice, Slovakia.
- Silvia Castro and Andreas König and Helwig Löffelmann and Eduard Gröller. **Transfer Function Specification for the Visualization of Medical Data.** *Technical Report* at the Institute of Computer Graphics and Algorithms, Vienna University of Technology, TR-186-2-98-12, 1998.

7.1 Introduction

When employing direct-volume-rendering for the visualization of volumetric data sets, typically a transfer function is used to classify the data. Data values are mapped into optical properties like color or transparency information. During the rendering procedure the transfer function is evaluated for specific data values in order to gather the contribution of certain voxels or sampled locations to the resulting visualization.

In general a transfer function τ is defined as a mapping from a cartesian product of scalar fields F to a cartesian product of optical properties O :

$$\tau : F_1 \times F_2 \times \dots \times F_n \rightarrow O_1 \times O_2 \times \dots \times O_m$$

The values of the dimensions n and m usually have to be kept small due to the overwhelming specification effort by the user. Typically a transfer function maps density values ($n = 1$) to opacity and color properties ($m = 2$). All other parameters which are necessary for shading the rendered voxels are determined by a more or less sophisticated illumination model [Max95]. Transfer functions of higher complexity usually take the magnitude of the gradient within the volume-data set into account ($n = 2$) [Lev88]. Even for these restricted cases ($n, m \leq 2$) the effort for specifying the transfer function is not to be neglected in daily clinical practice. Especially medical doctors with little experience in

computers or the mathematical background of volume rendering are usually not able to handle complex paradigms for specifying high-dimensional functions needed for the visualization.

Several approaches to provide a more user-friendly methodology for the specification of transfer functions have been proposed. Some methods analyze input data [FAT99, KD98] or output images [FBT98] to provide the user with an initial transfer-functions setup. When intuitive interaction mechanisms are missing, it is still hard for the user to fine-tune these proposed transfer functions to his needs. Other approaches generate an initial set of transfer functions, which is evolved for some generations by a genetic algorithms scheme [HHKP96]. It was also proposed to provide the user with an arranged display of previews, covering as much as possible of the entire domain of the transfer-function specification-parameters [MAB⁺97]. This approach proved to be very useful, as the user is just required to choose an appropriate region in the space of transfer function dimensions in order to gain an initial set of parameters for fine-tuning. The major drawback of this approach was the amount of preprocessing time needed for rendering the large number of previews covering the entire transfer function domain. With the rise of rendering hardware dedicated to generating direct-volume-rendering displays of large data sets with interactive frame rates [PHK⁺99], this problem can be overcome. In this chapter a strategy similar to this approach is presented. With the possibility of high speed rendering, even more properties in the possible space of transfer-function input-dimensions can be explored. The goal of the developed approach is to provide medical doctors with a system, which can be used intuitively. Nevertheless the possibilities for the specification of transfer functions must not be restricted.

The following section describes the approach presented in this chapter. Then the tools for the specification of contributions from the domains of data value, color, and opacity are described. The results section describes a prototype-like software implementation of the proposed system. Eventually conclusions are stated.

7.2 Advanced transfer-function specification

The specification of transfer functions usually is a difficult task. Appropriate settings have to be found in several search domains of sometimes infinite resolution. Trying to find a useful transfer function by moving a single sample through this high-dimensional search space is not possible. The approach presented in this chapter tries to define the specification step by step. If it is possible, the search-space domains are treated separately. Appropriate settings are found independently for these domains. The domains used as examples for the demonstration of the presented approach are:

- data range (center, width, and shape of contribution peaks)
- color (hue, luminosity, saturation)
- opacity (for different contributions to the visualization)

The presented approach is easily extendable by using different or additional specification domains, like gradient magnitude, principal curvature magnitude [HKG00], etc.

For each specification domain, different user interaction approaches have been developed. A different approach will be chosen by the user, depending on the experience of the user, a-priori knowledge about the investigated volume data set, the dedication to fine-tuning, the need for efficient generation of the transfer function in time-critical fields of application, and the suitability of predefined settings.

The basic concept of our approach is shown in figure 7.1. Contributions to the visualization are composed by choosing a certain number of intervals in the domain of the data values (“*peaks*”), selecting colors for these ranges and having them combined with specified opacities into a single transfer function. If i contributions are to be used, then the transfer function $f = \{f_o, f_c\}$ mapping data values d resolves to:

$$f_o(d) = \sum_i p_i(d) * o_i$$

$$f_c(d) = \begin{cases} c_i & p_i(d) \neq 0 \\ \text{black} & \text{else} \end{cases}$$

where $f_o(d)$ is the opacity definition of the transfer function to be specified, $f_c(d)$ the color definition of the transfer function, $p_i(d)$ is the contribution of a certain region of the data domain, o_i the opacity combination values for the contributions, and c_i are the colors for single contributions. When peaks $p_i(d)$ are overlapping, colors c_i have to be interpolated. An appropriate color space (like HLS) has to be used for the interpolation.

The core component of our specification scheme is a visualization display capable of rendering images of the investigated volume data set with real-time frame rates (see figure 7.2). Utilizing the VolumePro hardware ray-casting system [PHK⁺99], this rendering performance can be achieved. Whenever the user is specifying a contribution from the different domains, the influence of the changes made to the 3D rendering are displayed in this window. Note, that not only the final transfer function will be used for the display, but also transfer functions composed by the partially defined contributions from the domains (like certain data only, colored data regions only, and finally the composed transfer function). Small preview images rendering the volume data set using the intermediate (or final) transfer-function components are generated. Clicking one of these preview images, the user will activate the transfer function used for the preview. The rendering window will display the volume-data set classified with this transfer function, allowing the user to do further investigations by rotating and zooming the view. Using this approach, every manipulation of the transfer function by the user will result in immediate feedback of the changes in the 3D-rendering display (unlike other approaches, which try to give real-time feedback by classifying and coloring 2D slices).

7.3 Data range

The first step in the composition of a visualization is the specification of certain data ranges $p_i(d)$, which shall contribute to the visualization (e.g., skin, bone, etc). In the current prototype-like implementation, the contributions are specified in the domain of the data values of the volume data set. Nevertheless,

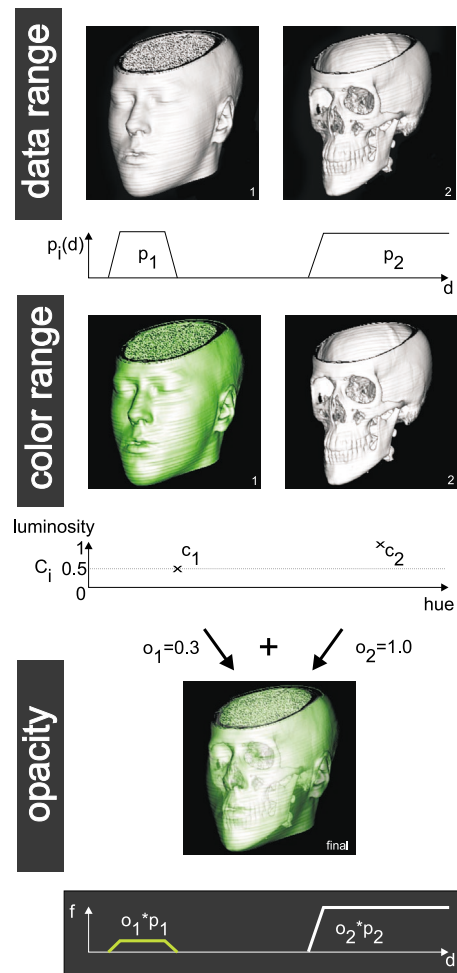


Figure 7.1: Two contributions are combined for the visualization. Parameters for the dimensions of the transfer function (data range, color, and opacity) are defined independently.

any scalar dimension of the data could be used for the selection, e.g., gradient magnitude, principal curvature magnitude, etc.

The contributing regions of the data domain are called *peaks* (see figure 7.3). The main properties of a peak are its location (*peak center* c) and the range of its contribution (*peak width* w). The contribution characteristic of specific data values within the range of a peak is defined by the shape of the peak. The following shapes have been proven to be useful:

- *Trapezoid:* The most common peak shape used with transfer functions is the trapezoid. Data values covered by the constant plateau around the center are mapped to full opacity, whereas the slopes are mapping data values to linearly increasing or decreasing opacity values. The width of the trapezoids slopes is an additional peak property definable by the user.




Figure 7.2: Real-time feedback for all modifications by the user is provided by a rendering window, using VolumePro technology for the calculation of the direct volume rendered image.

The definition of a trapezoid-peak shape is as follows:

$$f(t) = \begin{cases} 0 & 0 \leq t \leq c - \frac{w}{2} \\ \frac{t - (c - \frac{w}{2})}{s} & c - \frac{w}{2} < t \leq c - \frac{w}{2} + s \\ 1 & c - \frac{w}{2} + s < t < c + \frac{w}{2} - s \\ 1 - \frac{t - (c + \frac{w}{2} - s)}{s} & c + \frac{w}{2} - s < t \leq c + \frac{w}{2} \\ 0 & c + \frac{w}{2} < t \leq 1 \end{cases}$$

where c is the center of the trapezoid, w is the width (including the slopes), and s is the width of the trapezoids slopes.

There are two variations of the trapezoid-shaped peak, which have to be taken special care of: whereas the *tent* ($s = w/2$) is useful to extract information which is represented in a very narrow region within the data volume, the *box* shape ($s = 0$) should be avoided, as it usually produces artifacts due to the discontinuities at the vanishing slopes. Figure 7.3 gives a comparison.

- *Ramp*: Radiologists have been using the metaphor of *windowing* in 2D PACS software for years. A simple grey-level ramp (with certain center and width properties) is used to map data values to luminosity values. Data values below the minimal extend of the window are mapped to black (zero intensity), whereas values above the range of influence of the window are mapped to white (full intensity). The users are moving this “window” through the range of data values in order to enhance the visual contrast for certain regions of the data domain. As this is a well established technique in 2D, it was desirable to adapt this approach for 3D. Using a ramp-shaped peak, 2D windowing can be simulated in 3D. As we are utilizing VolumePro technology, the same interactive frame-rates as for 2D-windowing are

achievable. The definition of a ramp-peak shape is as follows:

$$f(t) = \begin{cases} 0 & 0 \leq t \leq c - \frac{w}{2} \\ \frac{t - (c - \frac{w}{2})}{w} & c - \frac{w}{2} < t \leq c + \frac{w}{2} \\ 1 & c + \frac{w}{2} < t \leq 1 \end{cases}$$

Figure 7.3 gives a comparison of the influences of these peak shapes.

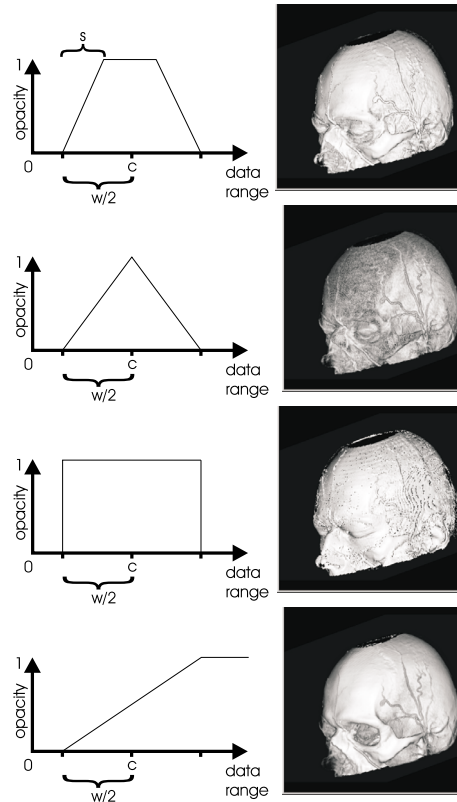


Figure 7.3: Different peak shapes (from top to bottom): trapezoid, tent, box, and ramp.

Several ways for specifying a peak have been included in the presented method. Depending on the task to be solved as well as on the level of experience of the user, a different approach may be chosen.

The central part of the user interface for this step is a bar displaying the range of data values (see figure 7.4). A grey-level ramp is displayed in this bar in order to give the user simple visual feedback which region of the data domain he is investigating.

When transfer function contributions (peaks) are displayed in the user interface, the numerical properties of a peak (center and width) are shown by plotted markers on this range bar, giving the user a hint, where the contribution he is investigating is located in the data domain.

By default the entire range of data values is displayed. Along with the bar, the limits of the currently active data range are given numerically (e.g., for a

data set with a signal resolution of 12 bit the number 0 would be displayed at the left end of the bar, whereas 4095 would mark the right-most extent of the bar). When the user wants to “zoom” into the data range (in order to get a new set of suggestions around an interesting region) the limits of the displayed range can be changed easily by activating a button and clicking the mouse at the position of the bar, where the new limit should be. Another button would be used to “reset” the limits of the bar and display the entire range of data values again.

A histogram of the data set is painted into the range bar. Peaks in this histogram give hints where certain materials are to be found in the data domain.

There are a number of possibilities to specify a peak. Depending on the a-priori knowledge about the data set, the experience of the user or the special demands of the visualization task, a different approach can be chosen. The following techniques have been proven to be useful:




Figure 7.4: Specification of contributions from the range of data values.

- *Accept a suggestion:* The system renders a certain number of preview images, using peaks covering the entire range of data values. Small thumbnail images are displayed (see figure 7.4). The user is free to activate one of the thumbnails by clicking on it and to investigate the visualization created with this transfer function peak in the rendering window. Due to the rendering speed of the VolumePro board, the visualization parameters (like camera position and orientation, lighting, shading, etc.) can be modified in real-time. If one of the suggested peaks produces, what the user is looking for the suggestion can be chosen as a result for this step.
- *Modify a suggestion:* If none of the suggestion shows exactly the features of the data set the user is interested in, the suggestion suiting the users needs best can be fine-tuned. Using his mouse the user drags one of the previews below the bar displaying the range of data values. The suggestion is turned into a modifiable peak. Its center can be moved through the range of data values by holding down the middle mouse button and moving the mouse horizontally. By moving the mouse vertically (again with holding the middle mouse button) the width of the peak can be modified. This approach has been chosen in analogy to the method of “2D-windowing” described above. Radiologists are used to select a certain range of the domain of data values and suppress the rest from rendering by holding the middle mouse button and dragging around the mouse. After a bit of training, this approach allows to fine-tune peak-parameters efficiently

without being aware of the actual numerical properties of the selection in the data domain which is currently rendered. This approach also allows to quickly browse through the entire data domain and therefore gain a fast overview on the features present in the volume. Without the real-time frame-rates provided by the VolumePro hardware, this approach would not be usable intuitively.

- *Use standard values:* For certain types of data sets, it is easily possible to directly derive the material scanned by the modality from the data value. For instance, for CT-data sets there is a linear mapping from data values to material densities. Certain materials are identifiable by their Hounsfield numbers. For this reason, predefined transfer function peaks can be used to display certain combinations of materials, which are interesting to the user. Common combinations include the display of bone, soft tissue, and contrast media (injected into vessels before scanning). The presented approach includes the possibility to load predefined sets of peaks and to save peak combinations the user wants to reuse. This approach is less suitable, when data sets acquired by MR or other non-xray-based modalities shall be visualized.
- *Numerical specification:* Another possibility is the specification of peak properties by their numerical representations. When for instance contrast medium of a certain manufacturer is used, its signal representation in the volumetric scan is known. The center and width values (and slope for trapezoids) for the peak rendering this feature may be entered manually by the user.

Whichever approach is chosen by the user, the goal of this specification step is one selected contribution from the range of data values for every feature, which has to be distinguished in the final visualization. Therefore this “data selection” step will be repeated by the user until the set of objects to be visualized is complete (and represented by thumbnail images in the user interface). The following step (described in the next section) yields similar selections from the domain of color values.

7.4 Color range

The next step in the specification process deals with the assignment of color information c_i to the ranges selected in the first step (see figure 7.5). For each peak defined in the first step, colors can be assigned independently. Color assignment is investigated for one “active” peak at a time, which is marked with a red frame in the user interface.

The main part of the user interface for color specification is a bar representing the hue and saturation domain of possible colors. Whenever a preview image is displayed in this user interface, it is connected to this color domain bar by a thin line, marking the color parameters for the color used to create the rendering. This enables the user to associate hue and saturation information with the presented color representation quickly. For the sake of convenience, the leftmost part of the color domain bar does not represent hue, but a grey-level ramp of entirely de-saturated colors with varying luminosity. Using this approach the

usually three-dimensional color space of a HLS-color system has been reduced to a two-dimensional specification scheme. As the quick specification of transfer function colors usually does not need other luminosity than an average value (0.5 is in the middle of full and no intensity in the HLS-color system), this approach seems to be sufficient. If the user wants to specify colors with other luminosities, other approaches can be used (see below).

Again a number of approaches are presented to the user:

- *Accept a suggestion:* The system renders a certain number of previews for the active peak, using colors covering the entire range of color hues. Small thumbnail images are displayed (see figure 7.5). The user is free to activate one of the thumbnails by clicking on it and investigate the visualization created with the active transfer function peak and the color represented by this preview in the rendering window.
- *Modify a suggestion:* If none of the suggestion shows exactly the color the user wants to see the investigated feature of the data set in, the color suggestion suiting the users needs best can be fine-tuned. Using his mouse the user drags one of the previews below the bar displaying the range of color values. By holding the middle mouse button and moving the mouse horizontally, the user changes the hue of the suggestion. Moving the mouse vertically changes the saturation values of the selected color. This approach has again been designed in analogy to 2D-windowing.
- *Specify hue or luminosity only:* If the user knows beforehand what color he is going to use for the peak in question, it is very convenient to “drag & drop” the peak-preview image right onto the color bar. If the user releases the mouse button in the leftmost part of the bar, a color-preview image in the grey-level range with a luminosity according to the mouse position is chosen. If the user drags to the right part of the color bar, the selected color is determined by the mouse position in the hue / saturation range. Luminosity is set to “middle”.
- *Choose a standard color:* By clicking a button, a set of standard colors is presented to the user. By choosing one of these, a color-preview image is created, which renders the volume data set using the selected peak (as opacity information) and the hue/luminosity/saturation value selected by the user.
- *Choose color parameters from graphical representation (color editor):* A standard color selection dialog is used to give the user the possibility to specify all parameters of a color (represented in different color spaces, like RGB or HLS) in an easy way. In the HLS-color space the user would for instance move the mouse around the hue/saturation circle and adjust the luminosity of the selected color by clicking into a region displayed with varying intensities. Figure 7.5 (lower right hand corner) shows the user-interface dialog used for this approach.
- *Numerical specification:* If the user is experienced with the use of different color systems, there is also the possibility to enter the characteristic values of a certain color numerically.




Figure 7.5: Specification of contributions from the range of colors.

The goal of this specification step is the assignment of color properties for every data-range peak selected in the first step. The final step will deal with the appropriate combination of these “colored peaks” in the opacity domain.

7.5 Opacity range

The final step in the specification process deals with the combination of the color-assigned data-range peaks prepared in the first steps into a single visualization. Different opacities o_i are assigned to the single peaks when they are combined into a single transfer function, which is used on the volume data set. The specification of these opacities o_i is the contribution of this step.

Again a user-friendly interface with real-time rendering feedback and a number of different approaches for the specification is provided (see figure 7.6).

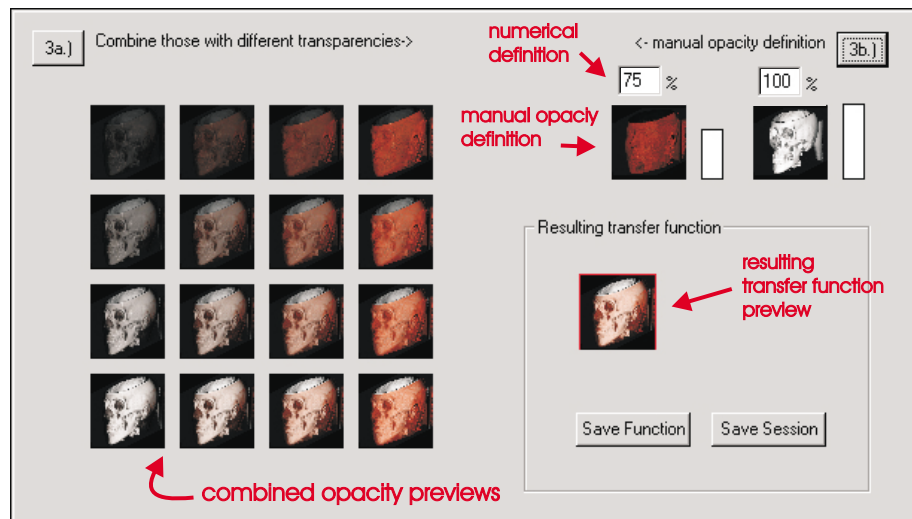


Figure 7.6: User interface for the combination of “colored peaks” by assigning opacities.

- *Choose suggestion:* The system is rendering a certain number of combinations of the “colored peaks” in discussion. When two peaks are to be combined, sixteen combinations with different amounts of opacity are displayed in a 4-by-4 matrix. The opacity contribution for the first peak is increasing in horizontal direction, whereas the opacity for the second peak is increasing vertically (see figure 7.6, left hand side). If three or more peaks are to be combined, more sophisticated schemes for the arrangement of the preview images have to be provided (e.g., sorted by dimensions).

The fastest way for the user to specify an appropriate combination is to simply select one of the preview images by clicking it with the mouse cursor.

- *Modify suggestion:* If none of the suggested combinations entirely satisfies the user, he may choose one of the preview images to have it put into “fine-tuning” mode. Modifications to the amounts of opacity used for the combination of the colored peaks can be made with the same interaction techniques as described below.
- *Manual definition:* The system displays the preview images for the prepared colored peaks along with small bars, representing the opacity contribution of the peaks (see figure 7.6, upper right hand corner). The user adapts the opacity contribution by simply clicking the middle mouse button into the bar and dragging the mouse in a vertical direction. Depending on the direction of the movement, the opacity contribution along with the graphical representation through the height of the bar is changed. Furthermore, the rendering window displays the changes to the final visualization in real-time during the modification of opacity contributions.
- *Numerical definition:* Finally, the user is able to enter the values for the contribution of the single colored peaks in numerical terms, giving the opacity as percentage of fully opaque.

The user usually modifies the opacities until the visualization rendered with the final transfer function fits his needs. Whenever settings defined in the first two steps seem not to be appropriate, the user would go back to data range or color specification and do modifications to the definitions there. Changes will of course be reflected in the definitions done in later steps.

7.6 Results

The concepts presented in this chapter have been implemented into the Tiani SpaceVision software system. A standard PC system (500MHz, 512MB RAM) was used as the hardware platform. It is equipped with a Real Time Visualization VP500 VolumePro board in order to provide interactive direct volume rendering performance. The prototype was developed using MS Visual C++ 6.0 running under Windows NT 4.0 SP5.

Figure 7.7 shows the entire user interface of the system. This screen shot has been taken after a user has defined three different contributions: skin, vessels (with contrast medium), and bone. These contributions have been assigned the colors green, red, and white. The user has chosen the opacity value of 24%

for the skin, 40% for the vessels, and 100% for the bone. The volume-data set rendered with the final transfer function is shown in the lower left hand corner.

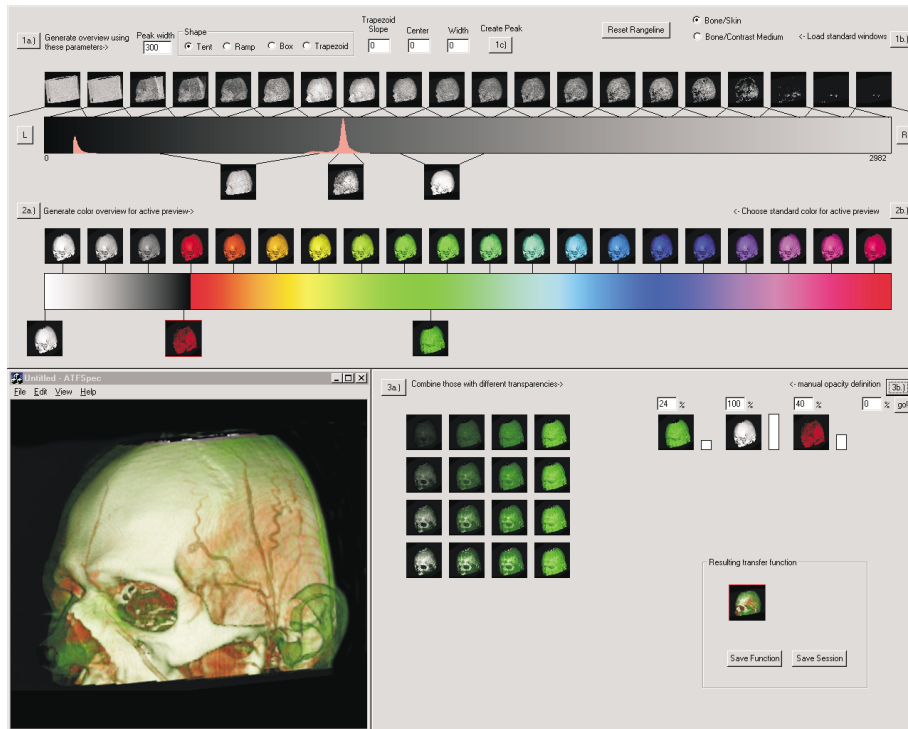


Figure 7.7: Specification of a transfer function with three contributions.

Other images created with the presented system can be viewed in higher resolution in the WWW: <http://www.cg.tuwien.ac.at/research/vis/vismed/ATFSpec>.

7.7 Conclusions

A new user interface paradigm for the specification of transfer functions in the field of medical volume visualization has been presented. The specification task is simplified by realizing the definition of mapping parameters for each search domain independently. A prototype software implementation (supporting data range, color, and opacity as search domains) has been used to prove the usefulness of the presented concepts. Several specification tools for each search domain are provided. Depending on the users experience and special demands of the visualization, different tools can be used. VolumePro technology has been used to provide the direct volume-rendering performance crucial for interactive feedback to the users when modifying the transfer function.

Conclusions

After this project on the applied visualization of medical data , i.e., after years of research in this field, quite some conclusions can be drawn. The most important insights we gained while building a medical 3D workstation include the following:

- The amount of available data is rapidly increasing. Soon there will be scanning modalities, which are capable of acquiring data sets of very high resolution within reasonable time frames. A lot of diseases will be detectable in early stages, thus dramatically increasing the chance to cure them.
- There are three main aspects, which have to be considered when building a medical workstation: usability, usability, and usability. Radiologists are no experts in the field of computers, let alone computer graphics.
- The quality of a medical workstation is measured by how simple it is to come up with incorrect findings. All techniques that might introduce artifacts or omit details, have to provide features, which allow the user to judge on the reliability of results.
- It is easy to change the way radiologists work, if the benefits and the quality of the offered tools are obvious. It is impossible, if this is not the case.
- The steady development of more powerful hardware and operating systems can be overwhelming. Nevertheless, it pays to be flexible and to permanently use the most powerful systems available.
- When techniques and approaches from the academic community are to be integrated into a commercial software system, a very tight feedback circle is crucial. The synergetic effects of a collaboration between visualization experts, software engineers, and medical professionals are invaluable, if a common basis for the communication is established.
- A lot of legal issues have to be taken into consideration in the high-risk area of medical applications. Reproducibility of results is a criterion, which is crucial to any software system in this field.
- Case studies and field-testing sessions are absolutely necessary. If a new approach works fine for twenty test-patient data-sets, it is not ensured, that there will not be any problems with the twenty-first. We found it useful to run test scenarios of up to eight months of daily clinical usage.

- A high degree of specialization is necessary. Just providing the user with a comprehensive toolbox of 3D functions will not yield an efficient usage. The application of 3D tools has to be focussed on the clinical tasks. There should be dedicated modules for specific tasks (e.g. investigation of coronary arteries, virtual bronchoscopy, etc), which integrate all the necessary functions and tools. An integration of a work-flow concept is also desirable. The system should guide the user through the usage of the different tools, providing an optimized order of application.

Summarizing, the development of a 3D medical workstation for the daily clinical use was quite a challenge. Nevertheless, fruitful collaborations have evolved, as radiologists are eager to get their hands on new tools, which enable them to do what is their job in a better and faster way: helping people.

Bibliography

- [AHH⁺94] R. Avila, T. He, L. Hong, A. Kaufman, H. Pfister, C. Silva, L. Sobierajski, and S. Wang. VolVis: A diversified volume visualization system. In R. Daniel Bergeron and Arie E. Kaufman, editors, *Proceedings of the Conference on Visualization*, pages 31–39, Los Alamitos, CA, USA, October 1994. IEEE Computer Society Press.
- [AN94] James Arvo and Kevin Novins. Iso-Contour volume rendering. In Arie Kaufman and Wolfgang Krueger, editors, *1994 Symposium on Volume Visualization*, pages 115–122. ACM SIGGRAPH, October 1994.
- [Beu91] S. Beucher. The watershed transformation applied to image segmentation. *Conf. on Signal and Image Processing in Microscopy and Microanalysis, Cambridge, UK*, September 1991.
- [BJNN97] Martin L. Brady, Kenneth Jung, HT Nguyen, and Thinkh Nguyen. Two-phase perspective ray casting for interactive volume navigation. In *Proceedings of the Conference on Visualization '97*, October 1997.
- [BL77] S. Beucher and C. Lantuejoul. Use of watersheds in contour detection. *International Workshop on Image Processing: Real-time Edge and Motion detection/estimation. Rennes, France*, September 1977.
- [BM97] W. Barrett and E. Mortensen. Interactive live-wire boundary extraction. *Medical Image Analysis*, (4):331–341, 1997.
- [BNKG00] Csebfalvi Balazs, L. Neumann, A. König, and E. Gröller. Gradient Estimation in Volume Data using 4D Linear Regression. *Proceedings of the Eurographics 2000*, 2000. Interlaken, Switzerland.
- [Bro92] L. G. Brown. A survey of image registration techniques. *ACM Computing Surveys*, 24(4):325–376, December 1992.
- [CRFS94] N. Coptly, S. Ranka, G. Fox, and R.V. Shankar. A data parallel algorithm for solving the region growing problem on the connection machine. *Journal of Parallel and Distributed Computing*, 21(1):160–168, April 1994.
- [CS92] D. W. Chakares and P. Schmalbrock. *Fundamentals of Magnetic Resonance Imaging*. Williams & Wilkins, 1992.

- [CS98] Wenli Cai and Georgios Sakas. Maximum intensity projection using splatting in sheared object space. *Computer Graphics Forum*, 17(3), September 1998.
- [Csé99] B. Csébfalvi. Fast volume rotation using binary shear-warp factorization. In E. Gröller, H. Löffelmann, and W. Ribarsky, editors, *Data Visualization '99*, Eurographics, pages 145–154. Springer-Verlag Wien, May 1999.
- [DCH88] Robert A. Drebin, Loren Carpenter, and Pat Hanrahan. Volume rendering. *Computer Graphics (SIGGRAPH '88 Proceedings)*, 22(4):65–74, August 1988.
- [Den90] Alan R. Dennis. An overview of rendering techniques. *Computers and Graphics*, 14(1):101–115, 1990.
- [DMMN99] K. K. Delibasis, G. K. Matsopoulos, N. A. Mouravliansky, and K. S. Nikita. Efficient implementation of the marching cubes algorithm for rendering medical data. *Lecture Notes in Computer Science*, 1593:989–1002, 1999.
- [EWE99] Klaus D. Engel, Rüdiger Westermann, and Thomas Ertl. Isosurface extraction techniques for web-based volume visualization. In David Ebert, Markus Gross, and Bernd Hamann, editors, *Proceedings of the 1999 IEEE Conference on Visualization (VIS-99)*, pages 139–146, N.Y., October 25–29 1999. ACM Press.
- [FAT99] Issei Fujishiro, Taeko Azuma, and Yuriko Takeshima. Automating transfer function design for comprehensible volume rendering based on 3D field topology analysis. In *Proceedings of the Conference on Visualization '99*, pages 467–470, 1999.
- [FBT98] Shiaofen Fang, Tom Biddlecom, and Mihran Tuceryan. Image-based transfer function design for data exploration in volume visualization. In *Proceedings of IEEE Visualization*, pages 319–326, 1998.
- [FKU77] H. Fuchs, Z. M. Kedem, and S. P. Uzelton. Optimal surface reconstruction from planar contours. *Commun. ACM*, 20:693–702, 1977.
- [FNVV98] A. Frangi, W. Niessen, K. Vincken, and M. Viergever. Multiscale vessel enhancement filtering. *Lecture Notes in Computer Science*, 1496:130–137, 1998.
- [Fri94] K. J. Friston. Statistical Parametric Mapping. In *Functional Neuroimaging*, pages 79–93, Academic Press, San Diego, CA, 1994.
- [FU00a] A. Falcao and J. Udupa. A 3D generalization of user-steered live-wire segmentation. *Medical Image Analysis*, (4):389–402, 2000.
- [FU00b] A. Falcao and J. Udupa. An ultra-fast user-steered image segmentation paradigm: Live wire on the fly. *IEEE Transactions on Medical Imaging*, 19(1), January 2000.

- [GHY98] R. Grzeszczuk, Ch. Henn, and R. Yagel. *Advanced Geometric Techniques for Ray Casting Volumes. SIGGRAPH 98 course #4 notes*. July 1998.
- [Gib98] S. F. F. Gibson. Constrained elastic surface nets: Generating smooth surfaces from binary segmented data. *Lecture Notes in Computer Science*, 1496:888–898, 1998.
- [Göp97] M. Göpfert. Segmentierung und Analyse von Gefäßbäumen in CT-Bildern zur präoperativen Einteilung der Leber in Segmente. *DKFZ, Heidelberg*, 1997.
- [HH92] K. Höhne and W. Hanson. Interactive 3D segmentation of MRI and CT volumes using morphological operations. *Journal of Computer Assisted Tomography*, 16(2):285–294, 1992.
- [HHKP96] Taosong He, Lichan Hong, Arie Kaufman, and Hanspeter Pfister. Generation of transfer functions with stochastic search techniques. In *Proceedings of IEEE Visualization*, pages 227–234, 1996.
- [HKG00] Jiri Hladuvka, Andreas König, and Eduard Gröller. Curvature-based transfer functions for direct volume rendering. *Spring Conference on Computer Graphics (SCCG) 2000*, 2000.
- [HMS95] Wolfgang Heidrich, Michael McCool, and John Stevens. Interactive maximum projection volume rendering. In Gregory M. Nielson and Deborah Silver, editors, *Proceedings of the Conference on Visualization '95*, pages 11–18. IEEE Computer Society Press, 1995.
- [HMSC00] Jian Huang, Klaus Mueller, Naeem Sharef, and Roger Crawfis. Fastplats: Optimized splatting on rectilinear grids. In *Proceedings of the Conference on Visualization (Vis) 2000*, 2000.
- [HO00] Rami Hietala and Jarkko Oikarinen. A visibility determination algorithm for interactive virtual endoscopy. In *Proceedings of the Conference on Visualization (Vis) 2000*, 2000.
- [HP96] K. H. Höhne and A. Pommert. *Volume Visualization. Brain Mapping: The Methods*, chapter 17. A. W. Toga and J. C. Mazziotta, Academic Press, San Diego, CA, 1996.
- [IK94] T. Itoh and K. Koyamada. Isosurface generation by using extrema graphs. In R. Daniel Bergeron and Arie E. Kaufman, editors, *Proceedings of the Conference on Visualization*, pages 77–84, Los Alamitos, CA, USA, October 1994. IEEE Computer Society Press.
- [Kan01] Armin Kanitsar. Advanced Visualization Techniques for Vessel Investigation. Master's thesis, Institute of Computer Graphics and Algorithms, Vienna University of Technology, January 2001.
- [KBD⁺98] Kevin Kreeger, Ingmar Bitter, Frank Dachille, Baoquan Chen, and Arie Kaufman. Adaptive perspective ray casting. In *IEEE Symposium on Volume Visualization*, pages 55–62, 1998.

- [KD98] Gordon Kindlmann and James W. Durkin. Semi-automatic generation of transfer functions for direct volume rendering. In *Proceedings of IEEE Volume Visualization*, pages 79–86, October 1998.
- [KKI99] Yukiko Kenmochi, Kazunori Kotani, and Atsushi Imiya. MARCHING CUBES METHOD WITH CONNECTIVITY. In *Proceedings of the 1999 International Conference on Image Processing (ICIP-99)*, pages 361–365, Los Alamitos, CA, October 24–28 1999. IEEE.
- [KSS98] Reinhard Klein, Andreas Schilling, and Wolfgang Straßer. Reconstruction and simplification of surfaces from contours. *Proceedings of the The Seventh Pacific Conference on Computer Graphics and Applications*, 1998.
- [LA98] Jianping Li and Pan Agathoklis. An efficiency enhanced isosurface generation algorithm for volume visualization. *The Visual Computer*, 13(9–10):391–400, 1998. ISSN 0178-2789.
- [Lac96] Philippe Lacroute. Analysis of a Parallel Volume Rendering System Based on the Shear-Warp Factorization. *IEEE Transactions on Visualization and Computer Graphics*, 2(3):218–231, September 1996.
- [LC87] WE. Lorensen and HE. Cline. Marching cubes: a high resolution 3D surface reconstruction algorithm. *Computer Graphics (SIGGRAPH '87 Proceedings)*, 21:163–169, 1987.
- [Lev88] Marc Levoy. Display of surfaces from volume data. *IEEE Computer Graphics and Applications*, 8(3):29–37, May 1988.
- [Lev90] Marc Levoy. A hybrid ray tracer for rendering polygon and volume data. *IEEE Computer Graphics and Applications*, 10(2):33–40, March 1990.
- [LG99] M. E. Leventon and S. F. F. Gibson. Model generation from multiple volumes using constrained elastic SurfaceNets. *Lecture Notes in Computer Science*, 1613:388–398, 1999.
- [LH91] David Laur and Pat Hanrahan. Hierarchical splatting: A progressive refinement algorithm for volume rendering. *Computer Graphics (SIGGRAPH '91 Proceedings)*, 25(4):285–288, July 1991.
- [LH98] Yarden Livnat and Charles Hansen. View dependent isosurface extraction. In *Proceedings of the 9th Annual IEEE Conference on Visualization (VIS-98)*, pages 175–180, New York, October 18–23 1998. ACM Press.
- [LL94] Philippe Lacroute and Marc Levoy. Fast volume rendering using a shear-warp factorization of the viewing transformation. In Andrew Glassner, editor, *Proceedings of SIGGRAPH '94 (Orlando, Florida, July 24–29, 1994)*, Computer Graphics Proceedings, Annual Conference Series, pages 451–458, July 1994.

- [LSJ96] Yarden Livnat, Han-Wei Shen, and Christopher R. Johnson. A Near Optimal Isosurface Extraction Algorithm Using the Span Space. *IEEE Transactions on Visualization and Computer Graphics*, 2(1):73–84, March 1996.
- [LY96] A. Law and R. Yagel. An optimal ray traversal scheme for visualizing colossal medical volumes. Technical report, Ohio State University, 1996. OSU-CISRC-2/96-TR11.
- [M⁺96] C. R. Mauer et al. Registration of 3-D Images Using Weighted Geometrical Features. *IEEE Transactions on Medical Imaging*, 15(6):836–849, December 1996.
- [MAB⁺97] J. Marks, B. Andalman, P. A. Beardsley, W. Freeman, S. Gibson, J. Hodgins, T. Kang, B. Mirtich, H. Pfister, W. Ruml, K. Ryall, J. Seims, and S. Shieber. Design galleries: A general approach to setting parameters for computer graphics and animation. In *Proceedings of ACM SIGGRAPH*, pages 389–400, 1997.
- [Mat94] S. V. Matveyev. Approximation of isosurface in the marching cube: Ambiguity problem. In R. Daniel Bergeron and Arie E. Kaufman, editors, *Proceedings of the Conference on Visualization*, pages 288–292, Los Alamitos, CA, USA, October 1994. IEEE Computer Society Press.
- [Max95] Nelson Max. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 1(2):99–108, June 1995. ISSN 1077-2626.
- [MB95] E. Mortensen and W. Barrett. Intelligent scissors for image composition. *Computer Graphics (SIGGRAPH '95), Los Angeles, CA*, pages 191–198, August 1995.
- [MCV⁺97] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens. Image Registration by Maximization of Mutual Information. *IEEE Transactions on Medical Imaging*, 16(2):187–198, April 1997.
- [Mey94] D. Meyers. Reconstruction of surfaces from planar contours, 1994. PhD thesis, University of Washington.
- [MGK99] L. Mroz, E. Gröller, and A. König. Real-time maximum intensity projection. In E. Gröller, H. Löffelmann, and W. Ribarsky, editors, *Data Visualization '99*, Eurographics, pages 135–144. Springer-Verlag Wien, May 1999.
- [MHG00] Lukas Mroz, Helwig Hauser, and Eduard Gröller. Interactive high-quality maximum intensity projection. *Computer Graphics Forum*, 19(3):341–350, August 2000.
- [MHS99] Michael Meißner, Ulrich Hoffmann, and Wolfgang Straßer. Enabling classification and shading for 3D texture mapping based volume rendering using OpenGL and extensions. In David Ebert, Markus Gross, and Bernd Hamann, editors, *Proceedings of the*

- Conference on Visualization '99*, pages 207–214, San Francisco, 1999. IEEE.
- [MKG00] Lukas Mroz, Andreas König, and Eduard Gröller. Maximum intensity projection at warp speed. *Computers and Graphics*, 24(3):343–352, June 2000.
- [MLG98] L. Mroz, H. Löffelmann, and E. Gröller. Bringing your visualization application to the internet. Technical Report TR-186-2-98-14, Institute of Computer Graphics, Vienna University of Technology, March 1998.
- [MMC99] Klaus Mueller, Torsten Möller, and Roger Crawfis. Splatting without the blur. In David Ebert, Markus Gross, and Bernd Hamann, editors, *Proceedings of the Conference on Visualization '99*, pages 363–370, San Francisco, 1999. IEEE.
- [MMS⁺98] Klaus Mueller, Torsten Möller, J. Edward Swan II, Roger Crawfis, Naeem Shareef, and Roni Yagel. Splatting errors and antialiasing. *IEEE Transactions on Visualization and Computer Graphics*, 4(2), April – June 1998. ISSN 1077-2626.
- [MSS92] David Meyers, Shelley Skinner, and Kenneth Sloan. Surfaces from contours. *ACM Transactions on Graphics*, 11(3):228–258, July 1992.
- [MSS94a] C. Montani, R. Scateni, and R. Scopigno. Discretized marching cubes. In R. Daniel Bergeron and Arie E. Kaufman, editors, *Proceedings of the Conference on Visualization*, pages 281–287, Los Alamitos, CA, USA, October 1994. IEEE Computer Society Press.
- [MSS94b] C. Montani, R. Scateni, and R. Scopigno. A modified look-up table for implicit disambiguation of Marching Cubes. *The Visual Computer*, 10(6):353–355, 1994.
- [MV98] J. B. A. Maintz and M. A. Viergever. A survey of medical image registration. *Medical Image Analysis*, 2(1):1–36, 1998.
- [MY96] Klaus Mueller and Roni Yagel. Fast perspective volume rendering with splatting by utilizing a ray-driven approach. In Roni Yagel and Gregory M. Nielson, editors, *Proceedings of the Conference on Visualization*, pages 65–72, Los Alamitos, October 27–November 1 1996. IEEE.
- [NDW93] J. Neider, T. Davis, and M. Woo. *OpenGL Programming Guide*. Addison-Wesley, Reading MA, 1993.
- [Neu00] Laszlo Neumann. Feature-preserving smoothing. Technical Report Technical Report TR-186-2-00-24, Institute of Computer Graphics and Algorithms, Vienna University of Technology, September 2000.
- [NH91] G. M. Nielson and B. Hamann. The Asymptotic Decider: Removing the Ambiguity in Marching Cubes. In *Proceedings of the Conference on Visualization '91*, pages 83–91, 1991.

- [NHM97] G. Nielson, H. Hagen, and H. Müller. *Scientific Visualization: Overviews, Methodologies, and Techniques*. IEEE Computer Society Press, 1997.
- [OM00] Jeff Orchard and Torsten Möller. Accelerated splatting using a 3D adjacency data structure. Technical Report TR 2000-05 School of Computing Science, Simon Fraser University, Burnaby, BC, Canada, November 2000.
- [OOLP98] Masato Ogata, TakaHide Ohkami, Hugh C. Lauer, and Hanspeter Pfister. A real-time volume rendering architecture using an adaptive resampling scheme for parallel and perspective projections. In *IEEE Symposium on Volume Visualization*, pages 31–38. IEEE, ACM SIGGRAPH, 1998.
- [PC92] Derek J. Paddon and Alan G. Chalmers. Parallel processing for rendering. *Third Eurographics Workshop on Rendering*, pages 1–8, May 1992.
- [PHK⁺99] H. Pfister, J. Hardenbergh, J. Knittel, H. Lauer, and L. Seiler. The VolumePro real time ray-casting system. *Proceedings of ACM SIGGRAPH*, pages 251–260, 1999.
- [RJ97] C. Revol and M. Jourlin. A new minimum variance region growing algorithm for image segmentation. *Pattern Recognition Letters*, (18):249–258, 1997.
- [RLF⁺98] K. Rehm, K. Lakshminaryan, S. Frutiger, K. A. Schaper, D. W. Summers, S. C. Strother, J. R. Anderson, and D. A. Rottenberg. A symbolic environment for visualizing activated foci in functional neuroimaging datasets. *Medical Image Analysis*, 2(3):215–226, 1998.
- [Rob93] Jonathan C. Roberts. An overview of rendering from volume data — including surface and volume rendering. Technical Report 13-93, University of Kent, Computing Laboratory, University of Kent, Canterbury, UK, December 1993.
- [RPSC99] Harvey Ray, Hanspeter Pfister, Deborah Silver, and Todd A. Cook. Ray casting architectures for volume visualization. In Hans Hagen, editor, *IEEE Transactions on Visualization and Computer Graphics*, volume 5 (3), pages 210–223. IEEE Computer Society, 1999.
- [RSA⁺94] K. Rehm, S. C. Strother, J. R. Anderson, K. A. Schaper, and D. A. Rottenberg. Display of merged multimodality brain images using interleaved pixels with independent color scales. *Journal of Nuclear Medicine*, 35:1815–1821, 1994.
- [SBTH92] T. Schiemann, M. Bomans, U. Tiede, and K. Höhne. Interactive 3D-segmentation. *Visualization in Biomedical Computing II, Proc. SPIE 1808*. Chapel Hill, NC, pages 376–383, 1992.

- [Sei99] Frank Seibert. Augmenting reality by using uncalibrated optical tracking. *Computers and Graphics*, 23(6):801–804, December 1999.
- [SFGS96] D. Schmalstieg, A. L. Fuhrmann, M. Gervautz, and Zs. Szalavári. 'studierstube' - an environment for collaboration in augmented reality'. *Proceedings of Collaborative Virtual Environments '96, Nottingham, UK, Sep. 19-20*, 1996.
- [SFYC96] R. Shekhar, E. Fayyadm, R. Yagel, and J. Cornhill. Octree-based decimation of marching cubes surfaces. *Proceedings of IEEE Visualization'96*, pages 335–342, 1996.
- [SGS95] Georgios Sakas, Marcus Grimm, and Alexandros Savopoulos. Optimized maximum intensity projection (MIP). In Patrick M. Hanrahan and Werner Purgathofer, editors, *Rendering Techniques '95*, Eurographics, pages 51–63. Springer-Verlag Wien New York, 1995.
- [Sha81] Michael Shantz. Surface definition for branching, contour-defined objects. *Computer Graphics*, 15(2):242–270, 1981.
- [Sho98] J. D. Shoeb. Improved Marching Cubes. <http://enuxsa.eas.asu.edu/shoeb/graphics/improved.html>, 1998.
- [SK00] M. Sramek and A. Kaufman. Fast ray-tracing of rectilinear volume data using distance transforms. In Hans Hagen, editor, *IEEE Transactions on Visualization and Computer Graphics*, volume 6 (3), pages 236–252. IEEE Computer Society, 2000.
- [SMM⁺97] J. Edward Swan II, Klaus Mueller, Torsten Möller, Naeem Shareef, Roger A. Crawfis, and Roni Yagel. An anti-aliasing technique for splatting. In Roni Yagel and Hans Hagen, editors, *Proceedings of the Conference on Visualization '97*, pages 197–204. IEEE, November 1997.
- [SNN95] R. Sammouda, N. Niki, and H. Nishitani. Neural network based segmentation of magnetic resonance images. *IEEE Nuclear Science Symposium*, 4:1827–1831, 1995.
- [SNS98] Y. Sato, S. Nakajima, and N. Shiraga. Three-dimensional multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. *Medical Image Analysis*, 2(2):143–168, 1998.
- [SS97] D. Schmalstieg and G. Schauffer. Smooth levels of detail. *Proceedings of VRAIS'97*, pages 12–19, 1997.
- [SWBN00] Y. Sato, C. Westin, A. Bhalerao, and S. Nakajima. Tissue classification based on 3D local intensity structures for volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 6(2), April 2000.
- [Tau95] G. Taubin. A signal processing approach to fair surface design. *Computer Graphics Proceedings*, pages 351–358, August 1995.

- [TD97] Klaus D. Toennies and Claus Derz. Volume rendering for interactive 3-D segmentation. In *Medical Imaging: Image Display*, volume 3031, pages 602–609. SPIE, 1997.
- [TT88] J. Talairach and P. Tournoux. *Co-planar Stereotaxic Atlas of the Human Brain*. Thieme Medical Publishers, Inc., New York, 1988.
- [vWHVP91] Theo van Walsum, Andrea J. S. Hin, Jack Versloot, and Frits H. Post. Efficient hybrid rendering of volume data and polygons. Technical Report DUT-TWI-91-17, Delft University of Technology, Department of Technical Mathematics and Informatics, Delft, The Netherlands, 1991.
- [WE97] R. Westermann and T. Ertl. A multiscale approach to integrated volume segmentation and rendering. *Computer Graphics Forum*, 16(3):117–128, August 1997. Proceedings of Eurographics '97. ISSN 1067-7055.
- [Wes90] Lee Westover. Footprint evaluation for volume rendering. In Forrest Baskett, editor, *Computer Graphics (SIGGRAPH '90 Proceedings)*, volume 24, pages 367–376, August 1990.
- [Wes91] Lee Alan Westover. SPLATTING: A parallel, feed-forward volume rendering algorithm. Technical Report TR91-029, Department of Computer Science, University of North Carolina - Chapel Hill, July 1 1991.
- [Wes95] Ruediger Westermann. Parallel volume rendering. In *Proceedings of the 9th International Symposium on Parallel Processing (IPPS'95)*, pages 693–699, Los Alamitos, CA, USA, April 1995. IEEE Computer Society Press.
- [WG92] Jane Wilhelms and Allen Van Gelder. Octree for faster isosurface generation. *ACM Transactions on Graphics*, 11(3):210–227, July 1992.
- [WK95] Sidney W. Wang and Arie E. Kaufman. Volume sculpting. In Pat Hanrahan and Jim Winget, editors, *1995 Symposium on Interactive 3D Graphics*, pages 151–156. ACM SIGGRAPH, April 1995.
- [WKB99] Ming Wan, Arie E. Kaufman, and Steve Bryson. High performance presence-accelerated ray casting. In David Ebert, Markus Gross, and Bernd Hamann, editors, *Proceedings of the Conference on Visualization '99*, pages 379–388, San Francisco, 1999. IEEE.
- [WNV00] O. Wink, W. Niessen, and M. Viergever. Fast delineation and visualization of vessels in 3-D angiographic images. *IEEE Transactions on Medical Imaging*, 19(4), April 2000.
- [WVH⁺00] Rainer Wegenkittl, Anna Vilanova, Balint Hegedüs, Daniel Wagner, Martin C. Freund, and Eduard Gröller. Mastering Interactive Virtual Bronchoscopy on a LowEnd PC. *Proceedings of the Conference on Visualization, Salt Lake City, USA, 2000*.

- [WVW94] Orion Wilson, Allen VanGelder, and Jane Wilhelms. Direct Volume Rendering via 3D Textures. Technical Report UCSC-CRL-94-19, University of California, Santa Cruz, Jack Baskin School of Engineering, June 1994.
- [XAW95] J. Xuan, T. Adali, and Y. Wang. Segmentation of magnetic resonance brain image: integrating region growing and edge detection. *2nd IEEE Intl. Conf. Image Processing, Washington, D. C.*, 1995.
- [YHW⁺97] Suyu You, Lichan Hong, Ming Wan, Kittiboon Junyaprasert, Arie Kaufman, Shigeru Muraki, Yong Zhou, Mark Wax, and Zhengrong Liang. Interactive volume rendering for virtual colonoscopy. In Roni Yagel and Hans Hagen, editors, *Proceedings of the Conference on Visualization '97*, pages 433–436. IEEE, 1997.
- [YK92] R. Yagel and A. Kaufman. Template-based volume viewing. *Computer Graphics Forum*, 11(3):153–167, September 1992. Proceedings of EUROGRAPHICS'92.
- [YS93] R. Yagel and Z. Shi. Accelerating volume animation by space-leaping. In *Proceedings of Visualization'93*, pages 62–69, October 1993.
- [Zen97] Marco Zens. Paralleles rendering. Master's thesis, Braunschweig University of Technology, Braunschweig, Germany, November 1997. Available from <http://diglib.eg.org/PubArc/diss>.
- [ZYS99] M. Nordin Zakaria, Md. Yazid, and Md. Saman. Hybrid shear-warp rendering. In *ICVC99 - International Conference on Visual Computing*, pages 144–153, 1999.

Disclaimer

The following designations are trademarks of TIANI Medgraph:

TIANI J-VISION/DIAGNOST
TIANI J-VISION/SPACEVISION

The following names and designations are the trademarks of other firms:

Windows	Microsoft Corporation
WinNT	Microsoft Corporation
Windows2000	Microsoft Corporation
UNIX	X/Open Company Limited
Java	Sun Microsystems, Inc.
VolumeZoom	Siemens AG
Somatom (Plus)	Siemens AG
Lightspeed	Phillips AG
Vitre	Vital Images
Advantage Workstation	General Electrics
Virtuoso	Siemens AG
VolumePro	Real Time Visualization
Plug'n'View	Voxar Ltd.

Other trademarks are the trademarks of their respective firms.

The patients' data used herein were made anonymous. The data sets used to create the images in this publication are courtesy of customers of Tiani Medgraph AG, Austria. Please check <http://www.tiani.com> for detailed information on Tiani Medgraph.

Related URLs

www.vismed.at – the official web-page of the project described in this thesis

www.tiani.com – Tiani Medgraph AG, the commercial host of the project described in this thesis

www.cg.tuwien.ac.at – Institute of Computer Graphics and Algorithms, Vienna University of Technology, the scientific host of the project described in this thesis

www.cg.tuwien.ac.at/research/vis/vismed – information, technical reports, and images in the scope of this project

www.cg.tuwien.ac.at/~koenig/PhD_thesis – provides this thesis in electronic form, along with high quality versions of the images, for downloading

Index

- 3D textures, 20
- abstract, i
- aluminium foam, 72
- aluminium foam cells, 75
- augmented reality, 57
- checkerboard, 37
- computed tomography (CT), 9
- contour connecting, 21
- CPR, 42
- CT angiography, 12
- curved planar reconstruction, 42
- cutting planes, 50
- data-set reformatation, 44
- DICOM, 32
- direct volume rendering, 46
- fMRI, 15
- focused views, 78
- frequency coding gradient field, 15
- fusion, 35
- future, 61
- hardware, 30
- hardware-accelerated rendering, 20
- head mounted display, 57
- high quality rendering, 54
- hybrid rendering, 53
- indexed volumes, 49, 66
- intelligent scissors, 22
- Larmor equation, 13
- live wire, 22
- magic mirrors, 54, 69
- Magnetic resonance imaging, 12
- marching cubes, 21
- mask smoothing, 21
- maximum intensity projection, 19, 51
- measurements, 45
- meister, 115
- MIP, 19
- mirrors, 54
- morphological operations, 22
- MPR, 42
- MR, 12
- MR angiography, 15
- Multi-Planar Reconstruction, 42
- multi-scale 3D-filtering based segmentation, 23
- multiple views, 68
- parallel volume rendering, 20
- PET, 16
- phase coding gradient field, 15
- precession, 13
- proton density weighted, 13
- quench valve, 15
- ray casting, 18
- reformatations, 40
- region growing, 22
- registration, 35
- render farm, 57
- rendering, 46
- scanning directions, 40
- screen-door, 37
- segmentation, 22, 39
- shape based interpolation, 34
- Shear Warp Factorization, 19
- shimming magnets, 15
- slice selection gradient field, 14
- sonography, 17
- SPECT, 17
- splatting, 19
- SSD, 21, 52
- surface rendering, 52
- Surface-based rendering, 21

- T1, T2, T*, 13
- thresholding, 22
- tools, 50
- tracking system, 57
- trademarks, 107
- transfer function specification, 83
- transfer functions, 48

- vessel tracking, 23
- virtual endoscopy, 23
- virtual reality, 57
- visualization, medical, 3
- volume filtering, 34
- volume of interest, 34
- volume reconstruction, 33
- volume rendering, 18
- volume sculpting, 34, 50
- VolumePro, 20

- watershed segmentation, 22
- wave algorithm, 23

- x-ray, 3

Curriculum Vitae



Personal information

Full name	Dipl.-Ing. Andreas Holger König
Born on	July 22th, 1974, in Graz, Austria.
Parents	Elfriede König (maiden name Perl) and Kuno König, running a company in the field of bibliopegistry.
Occupation	Vice President of Marketing, Tiani Medgraph AG, Brunn am Gebirge, Austria, www.tiani.com .
Contact information	Email: koenig@cg.tuwien.ac.at . WWW: http://www.cg.tuwien.ac.at/~koenig/ . Street address: A-7540 Neustift bei Güssing, Austria.

Education

- Volksschule (primary school) in Neustift bei Güssing, Austria, September 1980 - June 1984.
- Hauptschule (secondary school) in Güssing, Austria, September 1984 - June 1988.
- Bundesoberstufenrealgymnasium (high school) in Güssing, September 1988 - June 1992. Matura (final examination) with highest distinction on June 10, 1992.
- Studies of Informatik (computer science) at the Vienna University of Technology, October 1992 - April 1997. Diploma thesis "Real Time Simulation

and Visualization of NC Milling Processes for Inhomogeneous Materials on Low-End Graphics Hardware” at the Institute of Computer Graphics, up to October 1997. Graduation (with highest distinction) to “Diplom-Ingenieur der Informatik” (computer science) in October 1997.

- Further studies of Informatik (computer science) at the Vienna University of Technology, i.e., work on a PhD thesis on “Usability Issues in Volume Visualization” (part of the Vis Med project) at the Institute of Computer Graphics, since October 1997. Member of the scientific staff of the Institute of Computer Graphics, since October 1997.

Jobs

- Ferialpraxis (short job) at Alcatel Austria, Telecommunication Development, Austria; July 1995 - August 1995.
- System analyst and software developer at EDV GesmbH, Vienna. July 1996 - August 1997.
- Project Manager of the VisMed project, which was done as a cooperation of the Institute of Computer Graphics, Vienna University of Technology, and Tiani Medgraph GesmbH, Vienna. October 1997 - December 2000.
- Vice President of Marketing, Tiani Medgraph AG, Brunn am Gebirge, Austria, www.tiani.com, January 2001 - today.

Publications

- A. König. **3D Medical Visualization: Breaking the Limits of Diagnostics and Treatment.** ERCIM News (European Research Consortium for Informatics and Mathematics), number 3, January 2001.
- A. Vilanova i Bartrol, A. König, and E. Gröller: **Virtual Colon Flattening.** *3rd Joint EUROGRAPHICS - IEEE TCVG Symposium on Visualization (VisSym 2001)*, Ascona, Switzerland.
- Jiri Hladuvka, Andreas König, and Eduard Gröller. **Exploiting Eigenvalues of the Hessian Matrix for Volume Decimation.** *Winter School of Computer Graphics (WSCG 2001)*, Plzen, 2001.
- Csefalvi Balazs, L. Mroz, H. Hauser, A. König, and E. Gröller. **Interactive Non-Photorealistic Volume Rendering.** *22nd annual conf. of the European Association for Computer Graphics*, Manchester, UK.
- A. Vilanova i Bartroli, A. König, and E. Gröller: **Mastering Perspective Projection through Parallely Projected Slabs for Virtual Endoscopy.** *Spring Conference on Computer Graphics (SCCG 2000)*, Budmerice, Slovakia.
- Csefalvi Balazs, A. König, and E. Gröller. **Fast Surface Rendering of Volumetric Data.** *Winter School of Computer Graphics (WSCG 2000)*, Plzen, 2000.

- Csebfalvi Balazs, L. Neumann, A. König, and E. Gröller. **Gradient Estimation in Volume Data using 4D Linear Regression.** *Eurographics 2000*, Interlaken, Switzerland.
- Andreas H. König, Helmut Doleisch, Andreas Kottar, Brigitte Kriszt, Eduard Gröller. **AlVis – An Aluminium-Foam Visualization and Investigation Tool** *Joint Eurographics - IEEE TCVG Symposium on Visualization (VisSym00)*, Amsterdam, Netherlands.
- Jiri Hladuvka, Andreas König, and Eduard Gröller. **Curvature-based Transfer Function Specification.** *Spring Conference on Computer Graphics (SCCG 2000)*, Budmerice, Slovakia.
- H. Löffelmann, T. Theußl, A. König, and E. Gröller: **SMURF - a Smart sURFace model for advanced visualization techniques,** *Winter School of Computer Graphics (WSCG99)*, Plzen, 1999.
- L. Mroz, A. König, and E. Gröller: **Real-Time Maximum Intensity Projection,** In E. Gröller, H. Löffelmann, B. Ribarsky (eds.), *Data Visualization'99*, Springer Wien, pp. 135-144. Revised version "Maximum Intensity Projection at Warp Speed" published in the Journal Computers & Graphics, June 2000.
- A. Vilanova i Bartrol, A. König, and E. Gröller: **VirEn: A Virtual Endoscopy System,** Published in the *Journal Machine Graphics & Vision*, 1999.
- L. Mroz, A. König, and E. Gröller: **High Quality Maximum Intensity Projection.** Published in *Computer Graphics Forum*, 19(3), 1999.
- A. König, H. Doleisch, and E. Gröller: **Multiple Views and Magic Mirrors - Multimodal Volume Visualization of the Human Brain.** Published at the *Spring Conference on Computer Graphics (SCCG 99)*, Bratislava, Slovakia, 1999.
- Csebfalvi Balazs, A. König, and E. Gröller: **Fast Maximum Intensity Projection using Binary Shear-Warp Factorization,** *Winter School of Computer Graphics (WSCG 99)*, Plzen, 1999.
- S. Castro, A. König, H. Löffelmann, and E. Gröller: **Transfer Function Specification for the Visualization of Medical Data.** Technical Report at the Institute of Computer Graphics, Vienna University of Technology.
- Andreas König, and Eduard Gröller. **Real Time Simulation and Visualization of NC Milling Processes for Inhomogeneous Materials on Low-End Graphics Hardware.** Published at *Computer Graphics Interface (CGI '98)*, Hannover, Germany, June 1998.
- Helwig Löffelmann, Andreas König, and Eduard Gröller. **Fast Visualization of 2D Dynamical Systems using Virtual Ink Droplets.** Proceedings of the *Spring Conference on Computer Graphics (SCCG '97)*, Budmerice, Slovakia, pages 111-118, 1997.

Talks

- **3D Visualization.** Panel discussion at *Computer Aided Radiology and Surgery (CARS 2000)*, San Francisco, USA, 2000.
- **AlVis - Ein Werkzeug zur Analyse von Aluminiumschäumen.** *Symposium Metallschäume 2000*, Vienna, 2000.
- **AlVis – An Aluminium-Foam Visualization and Investigation Tool.** *Joint Eurographics - IEEE TCVG Symposium on Visualization (VisSym00)*, Amsterdam, Netherlands.
- **Das VisMed Projekt.** *University Management Club Vienna*, Austria, March 1999.
- **Real Time Simulation and Visualization of NC Milling Processes for Inhomogeneous Materials on Low-End Graphics Hardware.** *Computer Graphics Interface (CGI '98)*, Hannover, Germany, June 1998.
- **Issues Concerning the Usability of Volume Visualization.** *Scientific Learnshop on Visualization 98*, Bonn, Germany, February 1998.
- **Interleaved Dimension Decomposition: A New decomposition Method for Wavelets and its Application to Computer Graphics.** *Winter School on Computer Graphics (WSCG '98)*, Plzen, Czech Republic, February 1998.
- **Fast Visualization of 2D Dynamical Systems using Virtual Ink Droplets.** *Spring Conference on Computer Graphics (SCCG '97)*, Budmerice, Slovakia, April 1997.

Miscellaneous

- Chairman at the WSCG 98, Plzen, Czech Republic, February 1998.
- Reviewer for Visualization, SCCG, WSCG, VisSym, Computers and Graphics.
- Social Service (alternative military service) at AKH (general hospital) Vienna, Oct.1998 until Sept. 1999.

Hobbies

- Tae Kwon Do: I'm a beginner: 9. Cup
- Iaido: I've been searching perfection on the way of the sword for 8 years now. A lifetime to come...
- Archery: With compound / release / laser sight-style I'm on the high-tech side of this art.
- Motorbikes, survival training, reading and writing literature, rebuilding oltimers like Volkswagen Beetle or Porsche 356/911.

Andreas H. König, February 4th, 2001.

Acknowledgements

The last lines of this text should be words of gratitude to all those who helped to make this thesis possible. First of all, I have to thank **Meister Edi Gröller** for being a patient supervisor and for supporting this work with ideas, criticism, etc. I want to thank **Rainer Wegenkittl** for being a congenial colleague who I had the pleasure of discussing design issues with for the duration of this project. I am very thankful to **Helmut Doleisch**, **Andre Neubauer**, and **Armin Kanitsar**, who did great work in implementing the ideas presented in this thesis. For providing me with an environment and support for my work I want to thank **Werner Purgathofer**.

I'd like to thank all the VisMed team members, who contributed to this thesis: **Anna Vilanova i Bartroli**, **Balazs Csebfalvi**, **Jiri Hladuvka** (Institute of Computer Graphics and Algorithms), **Mario Bruckschwaiger**, **Egon Landschützer**, **Helmut Kopf**, **Gerd Höfler**, **Martin Capek**, and **Petr Felkl** (Tiani Medgraph).

Special acknowledgements go to **Tiani Medgraph AG**, Europe's leading PACS company, for supporting the project presented in this thesis by providing the financial resources and contacts to medical partners. **Franz** and **Martin Tiani** always filled this project with the spirit of technological excellence.

Furthermore, I would like to thank the following contributors from the field of healthcare and medical science:

- **Dr. Dominik Fleischmann**, the driving force behind the scientific collaborations we had with the General Hospital in Vienna.
- **Prof. Ewald Moser**, the fMRI guru at the General Hospital in Vienna.
- **Dr. Kai Kulenkampff**, Hospital Lainz, Vienna, who opened my eyes for the concerns of radiologists.

Lastly, but most importantly, I shall express my highest gratitude to my mother, **Elfriede König**. She supported me through the duration of my studies by providing the invaluable experiences of an entire lifetime. Without her, all this would not have been possible.